Integrated Photonics: Application-Specific Design and Manufacturing Series

Integrated Photonics for DATA COMMUNICATION APPLICATIONS

MADELEINE GLICK, LING LIAO
AND KATHARINE SCHMIDTKE

ELSEVIER

Integrated Photonics for Data Communication Applications

This page intentionally left blank

Integrated Photonics: Application-Specific Design and Manufacturing

Integrated Photonics for Data Communication Applications

Edited by

Madeleine Glick
Columbia University, New York, NY, United States

Ling Liao
Photonic Integration, Silicon Photonics Product Division, Intel Corporation, Santa Clara, CA, United States

Katharine Schmidtke
Meta Platforms Inc., Menlo Park, CA, United States

Series Editor

Lionel Kimerling
MIT Department of Materials Science and Engineering, Massachusetts Institute of Technology, Cambridge, MA, United States

Sajan Saini
MIT Initiative for Knowledge and Innovation in Manufacturing, Department of Materials Science and Engineering, Massachusetts Institute of Technology, Cambridge, MA, United States

ELSEVIER

Elsevier
Radarweg 29, PO Box 211, 1000 AE Amsterdam, Netherlands
The Boulevard, Langford Lane, Kidlington, Oxford OX5 1GB, United Kingdom
50 Hampshire Street, 5th Floor, Cambridge, MA 02139, United States

Copyright © 2023 Elsevier Inc. All rights reserved.

No part of this publication may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or any information storage and retrieval system, without permission in writing from the publisher. Details on how to seek permission, further information about the Publisher's permissions policies and our arrangements with organizations such as the Copyright Clearance Center and the Copyright Licensing Agency, can be found at our website: www.elsevier.com/permissions.

This book and the individual contributions contained in it are protected under copyright by the Publisher (other than as may be noted herein).

Notices

Knowledge and best practice in this field are constantly changing. As new research and experience broaden our understanding, changes in research methods, professional practices, or medical treatment may become necessary.

Practitioners and researchers must always rely on their own experience and knowledge in evaluating and using any information, methods, compounds, or experiments described herein. In using such information or methods they should be mindful of their own safety and the safety of others, including parties for whom they have a professional responsibility.

To the fullest extent of the law, neither the Publisher nor the authors, contributors, or editors, assume any liability for any injury and/or damage to persons or property as a matter of products liability, negligence or otherwise, or from any use or operation of any methods, products, instructions, or ideas contained in the material herein.

ISBN: 978-0-323-91224-2

For Information on all Elsevier publications
visit our website at https://www.elsevier.com/books-and-journals

Publisher: Matthew Deans
Acquisitions Editor: Stephen Jones
Editorial Project Manager: Zsereena Rose Mampusti
Production Project Manager: Erragounta Saibabu Rao
Cover Designer: Christian J. Bilbow

Typeset by MPS Limited, Chennai, India

Contents

List of contributors **xv**
About the editors **xxi**
Series foreword **xxiii**
Introduction **xxv**

1. Applications and key performance indicators for data communications **1**
Robert Stone, John Shalf, Doug Carmean, Ashkan Seyedi and Katharine Schmidtke
1.1 Introduction 1
1.2 Optical network case studies 1
1.2.1 High-performance computing 1
1.2.2 Enterprise data center 8
1.2.3 Hyperscale data centers—cloud computing 12
1.2.4 Hyperscale data centers—Web 2.0 13
1.3 Optical module form factors 19
1.3.1 Transceiver module architecture 21
1.4 Interconnect figures of merit 23
1.5 Major inflection points and challenges 24
1.5.1 Power and thermal 25
1.5.2 Resource disaggregation 25
1.5.3 Interconnect cost targets 28
1.5.4 Reliability 29
1.6 Considerations for future technology 30
References 31

2. Integrated lasers for data center silicon photonic-integrated circuits **35**
Chao Xiang and John E. Bowers
2.1 Introduction 35
2.1.1 Types of sources needed for datacom 36
2.1.2 Laser solutions 41
2.2 Integration issues 46
2.2.1 Process integration 46
2.2.2 Process temperatures 48
2.2.3 Yield 48
2.2.4 Reliability requirements 49
2.2.5 Cost 51

2.2.6 Coupling to Si waveguide 51
2.2.7 Optical stability and feedback tolerance 53
2.3 Performance requirements 53
2.3.1 Energy requirements 54
2.3.2 Threshold and slope efficiency 55
2.3.3 Output power 55
2.3.4 Spectral characteristics 56
2.3.5 Intensity noise 57
2.3.6 Operation temperature 57
2.4 State-of-the-art 58
2.4.1 Disaggregated 58
2.4.2 Integrated 58
2.5 Outlook 63
2.5.1 Next 5 years 63
2.5.2 Next 20 years 65
References 65

3. Optical modulators 69
Graham Reed, David Thomson, Weiwei Zhang, Frederic Gardes, Lorenzo Mastronardi, Ke Li, Shinji Matsuo, Shigeru Kanazawa, Laurent Vivien, Christian Lafforgue, John E. Bowers, Christian Koos, Marco Romagnoli, Marko Lončar, Mian Zhang, Stefan Abel and Ling Liao
3.1 Introduction 69
3.2 Modulation mechanisms 70
3.2.1 The thermo-optic effect 71
3.2.2 The free-carrier plasma dispersion effect 71
3.2.3 Electro-absorption modulation 73
3.2.4 Nonlinear optical effects 76
3.3 Historical review 77
3.4 Group IV modulators 80
3.4.1 Modulation methods 80
3.4.2 Examples of interferometric and resonant devices employing the Free-carrier plasma dispersion effect 82
3.4.3 Electro-absorption modulators 83
3.4.4 Modulator driver amplifiers 87
3.5 III−V modulators 93
3.5.1 Conventional III−V modulators 94
3.5.2 Membrane III−V modulators 96
3.6 Heterogeneous integration on silicon 97
3.6.1 III−V on silicon 97
3.6.2 Organic electro-optic materials on Si 99
3.6.3 Graphene on Si 102
3.6.4 Lithium niobate on Si 104

3.6.5 $BaTiO_3$ on Si 105
3.7 Summary and perspectives 108
References 109

4. High-speed photodetectors 123
Kazumi Wada, Jifeng Liu, Yasuhiko Ishikawa, Donghwan Ahn, Dong Pan, Pengfei Cai and Lionel Kimerling
4.1 Photodetectors in data center architecture 123
4.1.1 Performance parameters and waveguide−detector integration paradigm 124
4.1.2 Device structures 126
4.1.3 Ge versus III−V high-speed photodetectors 129
4.1.4 TxRx subsystem requirements 131
4.2 Epitaxy: dislocations and lattice strain in Ge/Si 132
4.2.1 Epitaxial growth of Ge on Si by chemical vapor deposition 132
4.2.2 Threading dislocation density in Ge 133
4.2.3 Tensile lattice strain by thermal stress 134
4.3 Advanced Ge photodiode processing on CMOS platform 136
4.3.1 Processing challenges with Ge photodetectors 137
4.3.2 Advanced integration of Ge devices with CMOS electronics 139
4.4 Avalanche photodetectors 142
4.4.1 Breakthrough and current market: Ge/Si heterojunction avalanche photodiodes in home and office communications 142
4.4.2 Avalanche photodiode principle and device performance of Ge/Si avalanche photodiode 143
4.4.3 Next-generation applications of avalanche photodiode 147
4.5 Photodetectors in data centers 149
References 149

5. Passive silicon photonic devices 159
Stefan Preble, Barton Bergman, Lewis G. Carpenter, Lukas Chrostowski, Amit Dikshit, Michael Fanto, Wenhua Lin, Matthew van Niekerk, Mohammad Rakib Uddin and Vijay Soorya Shunmuga Sundaram
5.1 Introduction 159
5.2 Survey of key passive optical devices 160
5.2.1 Waveguides 160
5.2.2 Routing elements 164
5.2.3 Splitters/combiners/couplers 168
5.2.4 Polarization diversity 170
5.2.5 Fiber-to-chip couplers 172
5.3 Multiplexers/demultiplexers and filters 175
5.3.1 Requirements and current trends 175
5.3.2 Interferometer filters and Mux/Demux 177

5.3.3 Resonator Mux/Demux 179
5.3.4 Waveguide Bragg grating Mux/Demux and filters 181
5.3.5 Arrayed waveguide grating and echelle diffraction gratings Mux/Demux 183
5.4 Future directions 185
5.4.1 Future directions in design and testing 186
5.4.2 Future directions in integration 186
References 187

6. Coherent interconnects for data centers 201
Clint Schow, Mark Filer and Christopher Doerr
6.1 Data center interconnect: system requirements 201
6.2 Direct detection and coherent links for DCI 203
6.2.1 Coherent optical devices 205
6.2.2 Evolution of DCI link technologies from intensity-modulated direct detection to coherent 206
6.3 Photonic-integrated circuits for coherent DCI 210
6.3.1 Photonic-integrated circuit material platforms 210
6.3.2 Coherent transceivers 210
6.3.3 Levels of electrical and optical integration 214
6.3.4 Performance details and limitations 214
6.3.5 Potential for further cost, size, and power reductions for coherent PICs 215
6.4 Development paths for volume deployment 219
6.4.1 The role of standards, consortia, and multi-source agreements 219
6.4.2 OIF 400ZR specifications 221
6.4.3 ZR + derivatives 221
6.4.4 Market impact of IPoDWDM shift 223
6.5 Future directions: outlook beyond 3 years 225
6.5.1 Interdatacenter interconnects 225
6.5.2 Intradatacenter interconnects 225
6.5.3 What is next in standards? 226
6.5.4 Summary 227
References 229

7. Photonic-integrated circuits for switched network interconnects 233
Saeed Fathololoumi, Yuliya Akulova, Ashkan Seyedi and David Hui
7.1 Data center interconnect: system requirements 233
7.1.1 Intradata center link requirements and standards 234
7.1.2 Electrical signal integrity 237
7.1.3 Photonic-integrated circuit loss bottlenecks and challenges 239
7.2 Integrated photonic building blocks for data center interconnect 241
7.2.1 Laser integration (external, flip-chip, heterogeneous) 242
7.2.2 Modulators 243

7.2.3 Receiver 244
7.2.4 Passives 244
7.3 Pluggable transceivers for intradata center interconnects 245
7.3.1 Directly modulated lasers 248
7.3.2 Externally modulated lasers 249
7.3.3 Si photonics platform 250
7.4 Photonic integration/copackaging for data center interconnects 254
7.5 Future directions 258
7.5.1 Scaling datacenters to 102-Tbps nodes and beyond 258
7.5.2 Photonic-integrated circuit scaling to higher data rates 259
References 260

8. Photonic switch fabrics in data center/high-performance computing networks 265
Qixiang Cheng, Chunhui Yao, Nicola Calabretta, Ripalta Stabile, Keijiro Suzuki, Hitoshi Kawashima, Weijie Tang, Madeleine Glick, Tao Chu, Kazuhiro Ikeda, Ryosuke Matsumoto, Shu Namiki, Keren Bergman and Richard Penty
8.1 Introduction 265
8.2 Key figures of merit for optical switching networks 266
8.2.1 System metrics 266
8.2.2 Device metrics 268
8.3 Overview of optical switch fabric architectures 270
8.3.1 Banyan-type blocking networks 270
8.3.2 Rearrangeable nonblocking networks 271
8.3.3 Wide-sense and strictly nonblocking networks 273
8.4 Overview of integrated optical switch technologies 274
8.5 Silicon photonic switch fabric 277
8.5.1 Thermo-optic Mach−Zehnder interferometer switch fabric 278
8.5.2 Electro-optic Mach−Zehnder interferometer switch fabric 280
8.6 InP optical switch fabric 281
8.6.1 All-active 16×16 semiconductor optical amplifier space switch 281
8.6.2 Active−passive 16×16 semiconductor optical amplifier space switch 281
8.6.3 Active−passive $8 \times 8 \times 8\lambda$ semiconductor optical amplifier space and wavelength switch 282
8.7 Heterogeneous photonic-integrated switch fabric 285
8.7.1 Gain integration with flip-chip bonding 285
8.7.2 Silicon-semiconductor optical amplifier gate wavelength selective switch 286
8.7.3 $Ge_2Sb_2Te_5$-on-silicon nonvolatile switch 287
8.8 Data center/high-performance computing optical interconnect architectures based on optical switches 288

8.8.1 Distributed deep learning training based on optical switching 288
8.8.2 Improved resource utilization through bandwidth steering enabled by optical switching 290
8.8.3 Disaggregated architecture based on nanoseconds optical switching 291
8.8.4 Highly scalable fast optical circuit switch networks 292
8.9 Remaining challenges and future perspectives 294
8.10 Summary 296
References 297

9. Photonic-integrated circuits and optical fabrics for heterogeneous computing systems 303
Doug Carmean, Robert Stone, Mark Wade, S.J. Ben Yoo, Ram Huggahalli and Larry Dennison
9.1 Problem statement 303
9.2 Applications space with requirements 305
9.3 Interfaces 308
9.4 System level reliability, resiliency, BER, lifetime 314
9.5 Power and cooling 318
9.6 Cost: current costs and path to $0.02/Gbps 320
9.7 Fabric architecture: control and data planes, scheduling 322
9.8 Photonic-integrated circuits for heterogeneous computing systems 324
9.8.1 Photonic-integrated circuits for photonic *I/O*s 324
9.8.2 Photonic-integrated circuits for enabling disaggregated reconfigurable heterogeneous computing systems 327
9.9 Conclusion 330
References 331

10. Photonic integrated circuit design methods and tools 335
Twan Korthorst, Wim Bogaerts, Duane Boning, Mitch Heins and Barton Bergman
10.1 Introduction 335
10.2 Photonic device design 338
10.2.1 Technology computer-aided design 338
10.2.2 Mode solvers 339
10.2.3 Propagation solvers 340
10.2.4 Electro/optical codesign 343
10.2.5 Feeding the circuit level design flow 343
10.3 Photonic-integrated circuit design 345
10.3.1 Formalisms 345
10.3.2 Process design kits and use of hierarchical design methodologies 346
10.3.3 Design flow: from idea to manufacturing 347
10.3.4 Schematic capture 348

10.3.5 Circuit simulation 349
10.3.6 Schematic driven layout 350
10.3.7 Back-annotation and postlayout simulation 351
10.3.8 Physical verification (design rule and layout vs. schematic checking) 351
10.3.9 Design for manufacturability, design for test, and post-layout data prep 352
10.3.10 Manufacturing, packaging, and test 352
10.3.11 Notes on electrical versus optical simulation and modeling 353
10.3.12 Process design kits 355
10.3.13 Process design kits for monolithic electro-optical processes 356
10.4 System design 357
10.5 Design for manufacturability 358
10.5.1 Manufacturing variations 358
10.5.2 Manufacturing variation-aware photonic models 359
10.5.3 Yield analysis and optimization 361
10.5.4 Photonics design rule checking, FILL, and optical proximity correction 363
10.5.5 Photonics design for test and design for reliability 364
10.6 Future photonic-integrated circuit design flow requirements 364
References 365

11. Photonic-integrated circuit fabrication and test approaches 369
Ajey P. Jacob, Sujith Chandran, Joris Van Campenhout, Marianna Pantouvaki, John Heck, Ken Giewont, Michal Rakowski, Yusheng Bian, Peter Debackere and Matthias Kuntz
11.1 Introduction 369
11.2 Relevant industry trends 369
11.3 Photonic-integrated circuit fabrication schemes and methodology 370
11.3.1 Photonic-integrated circuit physical design flow 371
11.3.2 Device design and simulation flow 372
11.3.3 Layout generation 372
11.3.4 Design rule check 373
11.3.5 Mask rule check 373
11.3.6 Design rule challenge: Manhattanization of curvilinear structures 374
11.4 Photonic-integrated circuit in-line and end-of-line testing 376
11.5 Photonic-integrated circuit integration, fabrication, and manufacturing 378
11.5.1 Hybrid photonics integration platform 378
11.5.2 Heterogeneous photonics integration platform 386
11.5.3 Monolithic photonics integration platform 391
11.5.4 InP III−V photonics integration platform 399

11.6 Photonic-integrated circuit outlook: fully monolithic photonic integration platform 406
References 407

12. Packaging and test technologies 411
Peter O'Brien, Kamil Gradkowski, Padraic E. Morrissey, Sylwester Latkowski, Helge Gehring, Wolfram Pernice and Jörg Smolenski
12.1 Introduction 411
12.2 PIC packaging technologies 412
12.3 Transceiver package design 412
12.4 Optical packaging 414
12.5 Electrical packaging 417
12.6 Hybrid laser-to-PIC packaging 420
12.7 Packaging design rules and standardization 420
12.8 Testing 421
12.8.1 PIC testing equipment 423
12.8.2 Standardized PIC testing protocols 425
12.8.3 The roadmap for PIC testing 426
12.9 3D printing of micro-optics for integrated photonics 427
12.9.1 Process for 3D printing of micro-optics 427
12.9.2 Application of 3D printing processes for micro-optics for integrated photonics 428
12.9.3 The roadmap for 3D printing of photonics 433
12.10 Conclusion 434
References 435

13. Reliability of Photonic-Integrated Circuits for data center and high-performance computing applications 439
Gloria E. Hoefler, Ranjani C. Muthiah, Randal A. Salvatore, Peter De Dobbelaere and Robert W. Herrick
13.1 Introduction 439
13.2 Reliability requirements for DC/HPC application environment 441
13.3 Reliability standards for DC and HPC applications 443
13.4 Methodology for PIC reliability 444
13.4.1 Failure mode effects analysis 444
13.4.2 Derisking 445
13.4.3 Design constraints 448
13.5 Subsystem and module assembly and integration 449
13.5.1 Reliability of photonics/electronics integration 450
13.5.2 Burn-in 452
13.5.3 Reliability of light source (not integrated in the PIC) 452
13.5.4 Reliability of lasers integrated into the PICs through heterogeneous integration 453
13.5.5 Reliability of fiber interfaces to PICs 455
13.5.6 Reliability of packaging for photonic subsystems 457

13.6 Reliability of DCI optical systems and field data 458
13.6.1 Module level reliability 459
13.7 Strategy for robust manufacturing quality 462
13.7.1 Wafer qualifications/certification 463
13.7.2 High volume manufacturing 463
13.7.3 On-going reliability monitoring (PIC and module) 464
13.8 Challenges and opportunities 466
13.8.1 Opportunities for new validation tools/technology 466
13.8.2 Challenges with emergent technologies 466
13.8.3 Opportunities across PIC capability and assembly improvements 467
References 467

Index **471**

This page intentionally left blank

List of contributors

Stefan Abel Lumiphase AG, Kilchberg, Switzerland

Donghwan Ahn School of Advanced Materials Engineering, College of Engineering, Kookmin University, Seoul, Korea

Yuliya Akulova Intel, Silicon Photonics Product Division, Santa Clara, CA, United States

S.J. Ben Yoo Department of Electrical and Computer Engineering, University of California Davis, Davis, CA, United States

Barton Bergman Dream Photonics, Vancouver, BC, Canada; American Institute for Manufacturing Integrated Photonics (AIM Photonics), Albany, NY, United States

Keren Bergman Columbia University, New York, NY, United States

Yusheng Bian GlobalFoundries, Malta, NY, United States

Wim Bogaerts Ghent University - imec, Ghent, Belgium

Duane Boning Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, Cambridge, MA, United States

John E. Bowers Department of Electrical and Computer Engineering, University of California, Santa Barbara, CA, United States

Pengfei Cai SiFotonics Technologies Co. Ltd, Woburrn, MA, United States

Nicola Calabretta Eindhoven University of Technology, Eindhoven, The Netherlands

Doug Carmean Meta Platforms Inc., Menlo Park, CA, United States

Lewis G. Carpenter American Institute for Manufacturing Integrated Photonics (AIM Photonics), Albany, NY, United States

Sujith Chandran Information Sciences Institute, University of Southern California, Marina Del Rey, CA, United States

Qixiang Cheng University of Cambridge, Cambridge, United Kingdom

Lukas Chrostowski The University of British Columbia, Vancouver, BC, Canada

Tao Chu Hangzhou, Zhejiang Province, Zhejiang University, P.R. China

Peter De Dobbelaere Cisco Systems Inc., Carlsbad, CA, United States

Peter Debackere Infinera, Sunnyvale, CA, United States

Larry Dennison NVIDIA Corporation, Santa Clara, CA, United States

Amit Dikshit American Institute for Manufacturing Integrated Photonics (AIM Photonics), Albany, NY, United States

Christopher Doerr Aloe Semiconductor, Inc., Middletown, NJ, United States

Michael Fanto Air Force Research Laboratory, Rome, NY, United States

Saeed Fathololoumi Intel, Silicon Photonics Product Division, Santa Clara, CA, United States

Mark Filer Omniva LLC, Cloud Technology, Santa Clara, CA, United States

Frederic Gardes Optoelectronics Research Centre, University of Southampton, Southampton, United Kingdom

Helge Gehring Westfälische Wilhelms-Universität Münster, Münster, Germany

Ken Giewont GlobalFoundries, Malta, NY, United States

Madeleine Glick Columbia University, New York, NY, United States

Kamil Gradkowski Tyndall National Institute, Cork, Ireland

John Heck Intel Corporation, Santa Clara, CA, United States

Mitch Heins Synopsys, Inc., Mountain View, CA, United States

Robert W. Herrick Intel, Santa Clara, CA, United States

Gloria E. Hoefler Infinera Inc., Sunnyvale, CA, United States

Ram Huggahalli Microsoft Corporation, One Microsoft Way, Redmond, WA, United States

David Hui Intel, Silicon Photonics Product Division, Santa Clara, CA, United States

Kazuhiro Ikeda The National Institute of Advanced Industrial Science and Technology, Tsukuba City, Japan

Yasuhiko Ishikawa Department of Electrical & Electronic Information Engineering, Toyohashi University of Technology, Tempaku Toyohashi, Japan

Ajey P. Jacob Information Sciences Institute, University of Southern California, Marina Del Rey, CA, United States

Shigeru Kanazawa NTT Device Technology Labs, NTT corporation, Atsugi-shi, Kanagawa, Japan

Hitoshi Kawashima The National Institute of Advanced Industrial Science and Technology, Tsukuba City, Japan

Lionel Kimerling MIT Department of Materials Science and Engineering, Massachusetts Institute of Technology, Cambridge, MA, United States

Christian Koos Institute of Photonics and Quantum Electronics (IPQ), Karlsruhe Institute of Technology (KIT), Engesserstr. 5, Karlsruhe, Germany

Twan Korthorst Synopsys, Inc., Mountain View, CA, United States

Matthias Kuntz Infinera, Sunnyvale, CA, United States

Christian Lafforgue Centre de Nanosciences et de Nanotechnologies, CNRS, Université Paris-Saclay, Palaiseau, France

Sylwester Latkowski Eindhoven University of Technology, Eindhoven, The Netherlands

Ke Li Optoelectronics Research Centre, University of Southampton, Southampton, United Kingdom

Ling Liao Photonic Integration, Silicon Photonics Product Division, Intel Corporation, Santa Clara, CA, United States

Wenhua Lin Intel, Santa Clara, CA, United States

Jifeng Liu Materials Science and Engineering Program Lead, Thayer School of Engineering, Dartmouth College, Hanover, NH, United States

Marko Lončar John A. Paulson School of Engineering and Applied Sciences, Harvard University, Cambridge, MA, United States

Lorenzo Mastronardi Optoelectronics Research Centre, University of Southampton, Southampton, United Kingdom

Ryosuke Matsumoto The National Institute of Advanced Industrial Science and Technology, Tsukuba City, Japan

Shinji Matsuo NTT Device Technology Labs, NTT corporation, Atsugi-shi, Kanagawa, Japan

Padraic E. Morrissey Tyndall National Institute, Cork, Ireland

Ranjani C. Muthiah Marvell Technology Inc., Santa Clara, CA, United States

Shu Namiki The National Institute of Advanced Industrial Science and Technology, Tsukuba City, Japan

Peter O'Brien Tyndall National Institute, Cork, Ireland

Dong Pan SiFotonics Technologies Co. Ltd, Woburrn, MA, United States

Marianna Pantouvaki imec, Leuven, Belgium

Richard Penty University of Cambridge, Cambridge, United Kingdom

Wolfram Pernice Westfälische Wilhelms-Universität Münster, Münster, Germany; Heidelberg University, Heidelberg, Germany

Stefan Preble Rochester Institute of Technology, Rochester, NY, United States

Michal Rakowski GlobalFoundries, Malta, NY, United States

Graham Reed Optoelectronics Research Centre, University of Southampton, Southampton, United Kingdom

Marco Romagnoli CNIT, Photonics Networks and Technologies Laboratory, Pisa, Italy

Randal A. Salvatore Infinera Inc., Sunnyvale, CA, United States

Katharine Schmidtke Meta Platforms Inc., Menlo Park, CA, United States

Clint Schow Department of Electrical and Computer Engineering, University of California Santa Barbara, Santa Barbara, CA, United States

Ashkan Seyedi NVIDIA Corporation, Santa Clara, CA, United States

John Shalf Lawrence Berkeley National Laboratory, Berkeley, CA, United States

Jörg Smolenski Nanoscribe GmbH & Co. KG, Eggenstein-Leopoldshafen, Germany

Ripalta Stabile Eindhoven University of Technology, Eindhoven, The Netherlands

Robert Stone Meta Platforms Inc., Menlo Park, CA, United States

Vijay Soorya Shunmuga Sundaram Rochester Institute of Technology, Rochester, NY, United States

Keijiro Suzuki The National Institute of Advanced Industrial Science and Technology, Tsukuba City, Japan

Weijie Tang Hangzhou, Zhejiang Province, Zhejiang University, P.R. China

David Thomson Optoelectronics Research Centre, University of Southampton, Southampton, United Kingdom

Mohammad Rakib Uddin American Institute for Manufacturing Integrated Photonics (AIM Photonics), Albany, NY, United States

Joris Van Campenhout imec, Leuven, Belgium

Matthew van Niekerk Rochester Institute of Technology, Rochester, NY, United States

Laurent Vivien Centre de Nanosciences et de Nanotechnologies, CNRS, Université Paris-Saclay, Palaiseau, France

Kazumi Wada Department of Materials Science and Engineering, Massachusetts Institute of Technology, Cambridge, MA, United States

Mark Wade Ayar Labs, Santa Clara, CA, United States

Chao Xiang Department of Electrical and Computer Engineering, University of California, Santa Barbara, CA, United States; Department of Electrical and Electronic Engineering, The University of Hong Kong, Hong Kong, China

Chunhui Yao University of Cambridge, Cambridge, United Kingdom

Mian Zhang John A. Paulson School of Engineering and Applied Sciences, Harvard University, Cambridge, MA, United States

Weiwei Zhang Optoelectronics Research Centre, University of Southampton, Southampton, United Kingdom

About the editors

Dr. Madeleine Glick is a senior research scientist at Columbia University where her current research focuses on applications of photonic devices and optical interconnects to bandwidth dense, energy-efficient computing systems. Madeleine received her PhD in physics from Columbia University in 1989. She subsequently conducted research on optical properties of III-V materials at the Department of Physics, Ecole Polytechnique Federale de Lausanne (EPFL) Lausanne, Switzerland. She was also a research associate with the European Organization for Nuclear Research, CERN, Geneva, Switzerland. In the early 2000s Madeleine initiated and led research activities on the use of photonics in computer networks. From 2002 to 2011, she was a principal engineer at Intel Research leading research on optical interconnects for data centers where she led one of the first research groups to explore optical interconnects and optical switching for computer networks and data centers. While at Intel, she initiated and led projects exploring digital signal processing for optical links—an impactful research collaboration with the University College London and Carnegie Mellon University. Her work has extended to research and collaborations with industry and universities, including the University of Oxford, University of Cambridge, MIT, and Lawrence Berkeley National Laboratory. Madeleine's contributions have been recognized in both the photonics and computing communities. She is one of the few optics researchers to have been published by top-tier computing conferences (Sigcomm 2016, SC19, SC20) and to be invited to the technical committee of the major supercomputing conference (SC18-SC22; vice chair of the Networks and Architectures subcommittee SC18). Madeleine is a fellow of Optica and the Institute of Physics UK (IOP). In 2022 she received the Distinguished Service Award from the IEEE Photonics Society.

Dr. Ling Liao is an Intel Fellow and chief architect of photonic integration in Intel's Silicon Photonics Product Division. She joined Intel in 1997 and spearheaded research in high-speed silicon modulation, optical transmitter integration, and copackaged optics. She currently leads the development of multi-terabit per second photonic engines to be copackaged with switch SOCs and XPUs for future power, cost, and bandwidth density scaling. Ling earned her BS and MS in materials science and engineering from the Massachusetts Institute of Technology and PhD in electrical engineering from the University of Surrey in the United Kingdom.

Dr. Katharine Schmidtke is the Director of ASICs and Custom Silicon Sourcing at Meta Platforms and plays a crucial role in developing the company's data center

optical technical strategy. With more than 25 years of experience in the photonics industry, she has held key positions at industry giants such as New Focus, JDS Uniphase, and Finisar in business development and product strategy. A fellow of Optica (formerly the OSA), Dr. Schmidtke is renowned for her commitment to open innovation and is a strong advocate of collaboration and ecosystem building in the optoelectronics industry. Katharine received her PhD in physics from the University of Southampton and conducted postdoctoral research at Stanford University.

Series foreword

The Information Age has been described as a technology progression from Communication and Computation to Imaging and Sensing to Autonomy and Learning. In each step of the way, fiber optic−led photonics, and now integrated photonics, has set the standard to usher data movement with unprecedented capacity, reliability, and fidelity.

In the last four decades, optical signal transmission revolutionized telecommunications, ultimately supporting a highly interconnected universe of four billion plus connections as of 2023, which continues to grow at a rate of 30% per year. High-performance computing reckons with an attendant computation appetite by scaling compute performance in Flop/s at a rate of $1000\times/10$ years. These developments have enabled universal information and data processing availability for machine learning, in addition to emergent augmented reality and artificial intelligence applications. A technology-driven community bringing together social, political, and economic objectives has emerged, and this global community will continue to rely on the continuation of data performance scaling.

Within the next decade, integrated photonics will continue to accelerate in providing ubiquitous enabling platforms for data flow and processing in ultralow power cloud computing and AI machine learning; ultrahigh bandwidth wireless networks (in tandem with radio over fiber data delivery); ultrasmart gas sensor and medical diagnostic lab-on-a-chips; ultrahigh resolution augmented reality telepresence and 3D imaging; quantum; and other novel computing paradigms.

The *secret sauce* for the future of information technology is imagination and innovation along the technology supply chain. This monograph series is intended to nurture that precious quality of creativity that built the present Age, with a consensus of visionary voices who strive to distill their cumulative insights, motivations, and prescriptions for building a photonic integrated circuit (PIC) technology responsibly and sustainably, into the manufacturing future. The groundswell of wisdom collected in this series' volumes is grounded in details and analysis, and at times purposefully polemic in its conclusions.

Each volume focuses on one application area: datacom, smart sensing, RF wireless, augmented imaging, quantum, and other computing or sensing paradigms. Each volume is guided by its focus area system requirements, to assess:

- the most critical application-specific design accomplishments of integrated photonics PDKs (process design kits);
- the most relevant performance metrics and benchmark values for near-term high-impact applications (electronic-photonic device, circuit, and system level);

- feasible tolerances for process variation during chip fabrication, and pragmatic design for manufacturing circuit layout practices with state-of-art process flows and fabrication tools; and
- design/fabrication standards for design for test, chip packaging and high-volume manufacturing, and efficient time-to-market product development cycles.

Guiding principles

The time is right for a progressive multivolume series that strives to establish a consensus-driven, systematic foundation for leading edge engineering in integrated photonics. With that in mind, the spirit of each volume encapsulates five guiding principles:

1. It is **inclusive** (the work will be accessible to cross-disciplinary engineering audiences).
2. It is **intercessive** (the work will prompt and mediate an elaboration of critical technology gaps in application areas).
3. It is **comparative** (the work will deeply address the intramural needs of an application area supply chain, and the work may briefly address intermural supply chain synergy with fellow application areas).
4. It is **conclusive** (the work will draw authoritative recommendations on necessary standards and practices for near-term design, fabrication, packaging, and testing).
5. It is **provocative** (the work will be a collective distillation of its many authorial voices, moderated by the editors with a bold yet sober recognition that integrated photonics remains a technology of still-uncharted potential).

It is our hope this series may serve as a formative technology library for a rapidly growing 21st-century integrated photonics workforce, one whose specialty training now includes an expert articulation of application-centric mass manufacturing and high-volume deployment needs for affiliated industry supply chains. Furthermore, the present-day converging and reformulating of multiple industry sectors signals the need for an increasingly cross-disciplinary and agile engineer, primed to potentially explore mastery-level training in more than one application area over the course of their career. We are at the threshold of a new electronics–photonics integration frontier, and it is the ambition of the many contributors to this series to help cultivate nothing less than a more *integrated form* of a modern PIC education.

We hope you will find the resulting book a worthwhile guide and reference, even as the integrated photonics industry continues to expand and to diversify, in the years to come.

Lionel Kimerling
MIT Department of Materials Science and Engineering,
Massachusetts Institute of Technology, Cambridge, MA, United States

Sajan Saini
MIT Initiative for Knowledge and Innovation in Manufacturing,
Department of Materials Science and Engineering, Massachusetts Institute of Technology,
Cambridge, MA, United States

Introduction

Welcome to the first volume in Elsevier's new monograph series on photonic integrated circuits. This volume provides a critical overview of the field of integrated photonics specific to high-capacity data communication in and between data centers and high-performance computing systems.

Integrated photonics is a rapidly growing technology that has the promise to transform the many application areas of photonics, similar to how integrated circuits transformed the electronics industry of the twentieth century. The most seminal application for integrated photonics is high-capacity data communication—where discrete photonics has played a major role in long distance communication for over half a century—but the technology will be challenged to satisfy the cost, density, and power needs of short-reach communication from data center interconnects to intrarack connectivity. Integrated photonics-based solutions are starting to address these challenges. They have progressively matured over the past decade and have seen volume deployment in data centers. This trend will continue and accelerate to meet the exponential growth of data traffic in and around data centers, and to enable new system architectures to scale performance.

The book starts by providing an overview of the data communication trends and key performance metrics of end users such as cloud computing companies, web-scale content providers, and government labs. The following chapters introduce each of the fundamental devices (lasers, modulators, photodetectors, and passive components) that are the elements in photonic circuits, reviewing device structure and design principles and their impact on function, as well as pivotal performance milestones. The subsequent chapters focus on bringing these devices together into photonic integrated circuits, designed to meet performance requirements of specific use-cases ranging from established applications connecting data networks over kilometers of optical fiber, to nascent applications interlinking chips inside computing equipment. The final chapters cover techniques and challenges critical to the commercial manufacturability of photonic integrated circuits, including the design tools and manufacturing processes being developed to support high-volume production, quality, and reliability.

The book is intended as a guide for engineers and researchers to provide a deeper insight into the devices and building blocks behind integrated photonics. It will appeal to readers curious to explore this rapidly evolving field and provide a more nuanced understanding of the present capabilities and the context for future innovation.

This page intentionally left blank

Applications and key performance indicators for data communications

1

Robert Stone[1], *John Shalf*[2], *Doug Carmean*[1], *Ashkan Seyedi*[3] *and Katharine Schmidtke*[1]
[1]Meta Platforms Inc., Menlo Park, CA, United States, [2]Lawrence Berkeley National Laboratory, Berkeley, CA, United States, [3]NVIDIA Corporation, Santa Clara, CA, United States

1.1 Introduction

Optical interconnects are an essential part of today's compute infrastructure, both within and between data centers (DCs) as well as enabling connectivity over transcontinental distances and to users. Continued exponential growth in traffic, doubling approximately every two years, [1] is placing sustained demands on technology evolution to match this growth rate in terms of both electronic and optical components. Given limitations on power delivery, cooling, and available space, this translates into a drive for higher efficiency and faster lane speeds, while reducing system power consumption. In this introductory chapter we introduce current state of the art applications, workloads and architectures drawing from DC, enterprise, and high-performance computing (HPC) case studies and identify key performance figures of merit and future challenges for optical interconnects.

1.2 Optical network case studies

1.2.1 High-performance computing

HPC interconnects have tended to accept more cost in the interconnect than the early days of hyperscale DCs. This is driven by the need to deliver high performance for scalable parallel applications where internode communication can become a dominant factor in usable application performance. In particular, the injection bandwidths for typical HPC systems have been much larger than for similar generation hyperscale systems. Sources of jitter such as contention or throughput bottlenecks can cause scaling of the performance of massively parallel application codes to asymptote at much lower concurrencies than the full scale of the machine. Cost sensitivities do exist for HPC interconnects (it cannot be unbounded), but the imperative to maintain scalable application performance limit how aggressively costs can be cut in this area.

Integrated Photonics for Data Communication Applications. DOI: https://doi.org/10.1016/B978-0-323-91224-2.00006-0
© 2023 Elsevier Inc. All rights reserved.

1.2.1.1 Workload requirements

The quality of an interconnect should be measured by how well it captures the communication requirements of a target application, as opposed to theoretical metrics such as diameter and bisection bandwidth, since such metrics depend only on the interconnect topology, ignoring the communication topologies of target applications. For this proposed approach, it is essential to understand scientific application communication requirements across a broad spectrum of computational methods. Once this information is derived, the interconnect design can be tailored for the specific communication requirements of the underlying applications in terms of cost and performance effectiveness, while exploring how new technologies can be adopted for breakthroughs in interconnect solutions.

HPC applications tend to have highly structured patterns that are persistent for long periods of time, whereas typical DC traffic has been considered mostly stochastic. However, the convergence of data analytics workloads and also irregular sparse computational patterns in HPC has created more diversity in those patterns. In order to quantify the requirements of HPC interconnect and study efficient implementation approaches, one must first develop an understanding of the communication characteristics for realistic scientific applications. Several studies have observed that many applications display communication topology requirements that are far less than the total connectivity provided by fully connected networks as shown in Fig. 1.1. For instance, the application study by Vetter et al. [2,3] indicates that the applications that scale most efficiently to large numbers of processors tend to depend on point-to-point communication patterns where each processor's average topological degree of communication is 3−7 distinct destinations, or neighbors. This provides strong evidence that many application communication topologies exercise a small fraction of the resources provided by fully connected networks.

However, it is important to understand that the structured nature of these communication patterns also leads to the potential for extreme hot-spots and huge disparities in link utilization. Random traffic patterns are an easier pattern because they distribute the load over many packet-switch buffers, whereas a bad mapping of the

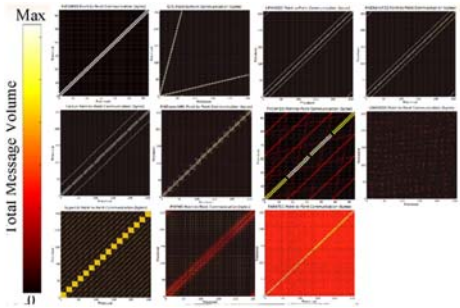

NAME	Discipline	Problem/Method	Structure
MADCAP	Cosmology	CMB Analysis	Dense Matrix
FVCAM	Climate Modeling	AGCM	3D Grid
CACTUS	Astrophysics	General Relativity	3D Grid
LBMHD	Plasma Physics	MHD	2D/3D Lattice
GTC	Magnetic Fusion	Vlasov-Poisson	Particle in Cell
PARATEC	Material Science	DFT	Fourier/Grid
SuperLU	Multi-Discipline	LU Factorization	Sparse Matrix
PMEMD	Life Sciences	Molecular Dynamics	Particle

Figure 1.1 Topological connectivity of a sampling of HPC applications, showing volume of communication at P = 256 for a cross section of HPC applications showing the structured and persistent nature of their communication patterns. *HPC*, High-performance computing.

structured application communication topology to the underlying interconnect topology can result in huge hot-spots. A number of remedies are employed to mitigate these hot-spots including (but not limited to):

1. aggressive forms of adaptive routing for HPC,
2. higher bisection bandwidths to reduce the impact of adversarial traffic as multiple applications must share the same system (albeit they do not typically share the same nodes), and
3. proactive congestion avoidance methods rather than reactive detection (e.g., AIMD for TCP protocol) or even the various forms of random early detection methods (RED) [4] that is typically used in commodity Ethernet.

Because many HPC applications are bulk-synchronous in nature, latency is also a significant consideration for the design of HPC interconnect topologies and protocols. When a bulk-synchronous application reaches a synchronization point, then the entire application must pause until the communication phase is completed. To the extent that latencies can be reduced either through more efficient communication protocols or packet encodings, or lower diameter topologies (such as Dragonfly —described in Section 2.3), latency and jitter can have a significant impact on application performance for this class of applications. More recently, advances in remote direct memory access (RDMA) and other means to better overlap communication with computation have emerged that have reduced the pressure on reducing interconnect latencies. With the propagation of light being limited to approximately 5 ns per meter over fiber, there is only so much further opportunity to reduce latencies in the long term.

1.2.1.2 Impact of emerging artificial intelligence workloads

AI has forced a convergence between hyperscale DC requirements and those of HPC as the injection bandwidths required to support scalable training on GPU-accelerated systems have grown to match that of typical HPC installations. One minor twist on conventional topologies is the HyperX, which combines a tree-type network for the core with a hypercube-like topology for the endpoints. This design choice is driven by the common ring-of-rings communication pattern for communicating the gradients for the weights for typical data-parallel machine learning (ML) training [5]. We also see more examples where the interconnection network at the endpoints switches to a proprietary higher bandwidth network such as NVLink/NVSwitch or AMDs xGMI where the proprietary memory fabric is used to create very fat nodes at the endpoints. This constitutes switching network link technologies and protocols at the leaf nodes of the network to support artificial intelligence (AI) application requirements rather than maintaining a uniform network protocol and link technology throughout the system.

More recently, the most rapidly growing area of ML has been recommender and transformer networks such as bidirectional encoder representations from transformers (BERT) that have millions or even billions of degrees of freedom—far too large to fit entirely within a single node, which requires a move from data-parallel to model-parallel in order to support the enormous memory footprint. This puts

even more extreme bandwidth requirements between GPUs, but also the sparsity of these networks increase the nonuniformity of load requirements between participating nodes. The exact impact on interconnect topologies is unclear at this point, but the growth and scale of the AI/ML market and the substantial profit margins mean that there is enough funding to support much more extreme, specialized designs at the leaf nodes to meet the needs of that targeted workload to the extent that they do not align with mainstream commodity link technologies and topologies. Scalable application performance is paramount and drives a departure from a purely cost-reduction standpoint in regard to interconnect technologies.

1.2.1.3 High-performance computing link technologies

Among the many issues affecting scalability of future system is the scaling of high bandwidth interconnects, designed for both on-chip and off-chip communication to memory and to other computational devices. Future computing systems, whether based on traditional circuits or the proposal nanotechnology devices, will rely on parallelism to keep power budgets manageable while increasing performance. At the macroscale, interconnects must keep costs, efficiency, and power consumption under control in the face of growth in system parallelism. These have been largely driven by the availability of link technologies over time. As mentioned earlier, the HPC applications have historically tended to have much higher injection bandwidth requirements than the typical web-farm and map-reduce-oriented architectures of the early enterprise/hyperscale DCs. The cost considerations drove non-HPC DCs toward commodity Ethernet, whereas the HPC tended to create more customized link technologies to serve its requirements. More recently, the rise of AI in hyperscale and enterprise data center (EDC) workloads has increased the demand for high injection bandwidth and systemwide throughput for those interconnects. This convergence of requirements has also led to a concomitant convergence of architectures.

Whereas cabling for HPC interconnects tended to be more proprietary in the early days of parallel computing, over time the link media has largely converged with hyperscale DCs, with both using active optical cables (AOCs) with standardized electrical interfaces (SFP, SFP + , etc.) in order to reduce costs by sharing with larger markets. Even for the full custom interconnects such as the Cray SlingShot, the links are standard Ethernet link technology even though the protocol running over those links is proprietary and targeted at HPC. This shows that although HPC is willing to spend more on the link technologies, there are limits. When it is not feasible to use customization to outperform the mainstream transceivers, then HPC has rapidly moved to adopting mainstream AOCs and connector technologies for the physical links while maintaining a custom protocol layer for the data that is transmitted over that fiber.

As the cost and complexity of moving data over copper will become more difficult over time, the cross-over point where optical technology becomes more cost-effective than electrical signaling has been edging closer to the board and chip package at a steady pace for the past two decades. Contemporary short-distance

copper links consume about 10−20 pJ/bit but could be improved to <1 pJ/bit for short-haul board-length links in the near future. However, the efficiency and/or data carrying capacity of the copper links will fall off rapidly with distances that forces movement toward optical links. Moreover, copackaged silicon photonics offers the promise of breaking through the limited bandwidth and packaging constraints of organic carriers using electrical pins. By becoming cotravelers with the recently injection-bandwidth-hungry hyperscale DC market, the HPC community is better able to share the R&D costs for these emerging photonic technologies.

1.2.1.4 High-performance computing topologies

In the early days of parallel computing, the higher injection bandwidth required for HPC nodes demanded much higher bandwidth link technologies such as HIPPI6400. The enormous link bandwidth made it challenging to build high radix switches, so the early topologies typically adopted topologies amenable to lower radix switch building blocks. Examples of these lower radix topologies include the 3D Torus or mesh networks, which were employed by successive generations of Cray XT series systems and the first two generations of IBM BlueGene systems. Another example of a topology driven by low-radix switch building blocks is the hypercube, which was employed by early parallel systems such as the "Cosmic Cube" and by the SGI Origin 2000 and successive generations of SGI DSMPs.

However, not all link technologies limited switching to lower radix. Higher radix switch technologies could be built up from smaller cointegrated building blocks with the most successful and widespread technologies included tree networks such as the Fat Tree and CLOS networks such as the interconnect on the Thinking Machines CM-5 and successive generations of Myrinet interconnected machines in the 1990s. Through the 2000s Ethernet and InfiniBand-based clusters favored the tapered-CLOS topology. Tapering of the CLOS for such interconnects tended to be less aggressive than for DC networks (3:1 or lower ratios) because of the need for high bisection. However, for nonadaptive source-routed networks, it would not be practical to use the full bisection at the spine switches due to unavoidable contention at the intermediate layers of the interconnect. This topology worked well for intermediate-sized clusters, but the very large systems struggled to scale as the number of additional layers required for the interconnect scale nonlinearly with the number of endpoints in the system.

High radix switches enabled not only more scalable hierarchical/tree networks but also development of new hierarchical all-to-all topologies such as the Dragonfly (employed in the Cray Aries and SlingShot systems) and the three-tier network topology employed in the BlueGene/Q that used a hybrid of a high-dimensional torus within the rack and all-to-all between racks. These network topologies enabled the copper-link technologies to be concentrated within the rack and high-speed optical links isolated to the all-to-all connectivity between the racks, which offered some cost savings at a time when optical links remained very expensive. With the all-to-all links in these topologies, they achieve extremely high bisection bandwidths. However, the aggressive adaptive routing in this class of interconnects

enabled dynamic trade-off of point-to-point bandwidth against the effective bisection bandwidth for the system. For example, if the point-to-point link between a pair of nodes becomes saturated, then the adaptive routing will then bounce data off of secondary and even tertiary nodes. Each bounce effectively reduces the effective bisection bandwidth but increases the effective point-to-point bandwidth between the affected nodes. This effect can cause some unpredictable variations in communication performance in response to adversarial traffic, but the scalability of component counts for very large systems and the reduction in the network diameter (thereby limiting the number of hops to traverse the interconnect) have made the Dragonfly popular in the largest systems in the Top500 [6].

1.2.1.5 Considerations for schedulers

Support for topologically agnostic greedy parallel schedulers is a consideration because scheduling efficiency can greatly impact net throughput of an HPC DC. Whereas typical hyperscale DCs operate at 45%−50% occupancy [7], a typical HPC DC will use batch queueing to operate at 95% occupancy on average throughout the year. This occupancy figure includes downtime for maintenance and node failures. A good scheduler will attempt to schedule jobs in topologically adjacent portions of the interconnect to minimize contention from adversarial traffic, but as jobs of different sizes get scheduled and retired, it can become increasingly difficult to bin-pack jobs into contiguous portions of the interconnect topology. Constrained or hierarchical topologies such as the torus and to a lesser extent the Fat Tree will see substantial performance degradation from such interjob contention (adversarial traffic) as job fragmentation increases. An example is shown in Fig. 1.2. If the scheduler preferentially schedules smaller jobs in subsections of the topology as they come available, it can starve the larger jobs of resources or unfairly delay their start time. Alternatively, a "fair-share" scheduler that gives equal time to small and large jobs

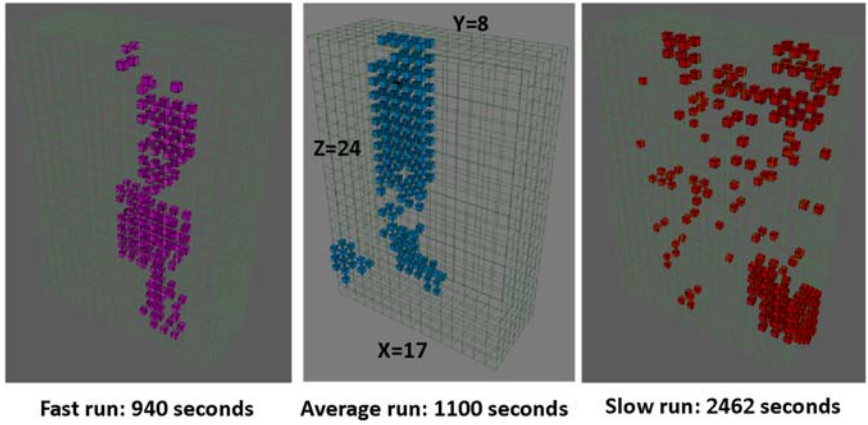

Figure 1.2 Job fragmentation for a 3D torus interconnect (Cray XT3) greatly impacts the delivered job performance for a climate modeling application (FVCAM).

can hold back on scheduling a job until a sufficiently large contiguous region of the interconnect comes available, but this can significantly impact system utilization as it would reduce occupancy as the scheduler must leave nodes idle as it waits for large enough contiguous sections. Therefore flatter interconnects, adaptive routing methods, or all-to-all topologies such as Dragonfly can offer a friendlier environment for schedulers to maximize system utilization through high occupancy.

1.2.1.6 Historical trends in high-performance computing technology usage from the Top500

For as long as the Top500 list has been in existence, HPC system measured LINPACK performance that has been consistently increased by a factor of 1000x every 11 years. Although the Top500 has been much maligned given the FLOP-heavy characteristics of its LINPACK benchmark and is no longer considered representative of modern application requirements, the accumulated data over time of the systems on that list offer a treasure trove of information about trends in HPC system architectures over the past 25 + years of its existence. Moore's law has supplied 100 × of that improvement, but the extra 10 × has been delivered through innovations that are specific to the leading-edge HPC architectural space. In particular, the burden of this extra 10 × scaling has fallen largely on advanced interconnect architectures that enable large arrays of commodity components to function together as an integrated, scalable HPC system.

Ethernet is used by many parallel HPC systems on the Top500 list throughout its history, but since the introduction of InfiniBand (IB), a majority of systems use IB as the preferred fabric. Although the LINPACK benchmark is well known to be FLOP-heavy, it also happens to be extremely challenging for the interconnect—especially at the end of the benchmark where the panel sizes get smaller and communication becomes more frequent and intense. Whereas R_{max} is the peak theoretical speed of a system, the R_{peak} is the top measured performance, and the difference between R_{max} and R_{peak} is informative as it shows the net efficiency of the system. Whereas the typical Ethernet-based clusters achieve a 50%−70% of peak performance, the efficiency of InfiniBand systems has typically been much higher. However, the emergence of HPC-Ethernet and the convergence of physical layers (PHYs) may change this balance over time. But the dominance of InfiniBand thus far does appear to be correlated with evidence of superior efficiency delivered for communication-intensive workloads. The recent emergence of HPC-Ethernet, with more advanced congestion avoidance features, has greatly improved the efficiency of Ethernet-based systems in the list as shown in Fig. 1.3, and so the advantages offered by InfiniBand have begun to fade.

Although commodity interconnects such as Ethernet and InfiniBand dominate the Top500 list, they are primarily dominant in the lower 490 systems. When we look at the Top10 systems on the list, we tend to see more custom/proprietary interconnect designs such as the Fujitsu TOFU, and the Cray SeaStar, and Aries. As described earlier, the congestion avoidance features required for the 2020 Top500 list show how the proprietary category of interconnects dominates in the largest systems with

a handful of smaller copies of those same systems filling out the interstices. This is depicted in Fig. 1.4. So overall, the bottom 90% of the Top500 has many commodity interconnect systems, whereas the top-tier systems tend to use more proprietary interconnect technologies. In the space of commodity interconnects, IB tended to dominate over Ethernet until the advent of more effective HPC-Ethernet protocols, which enabled more efficient performance and scalability for HPC applications. In the commodity space, the performance difference between Ethernet and InfiniBand has become more or less indistinguishable in the most recent lists.

1.2.2 Enterprise data center

Modern EDCs are carefully designed so that they can optimize for interconnect cost efficiency, enable network robustness, and minimize time to a solution for a given workload. This can initially be enabled by keeping the interconnect cable length below 50 m, where most of the links are 10 m and shorter. This allows copper cables, actives or passive, to dominate. Secondarily, by using a topology with redundant paths, one can ensure operation through hardware or infrastructure faults. As a result, network topologies that favor short interconnect lengths such as Dragonfly (DF) shown in Fig. 1.5A [8] and enable robust operation such as Fat Tree (FT) shown in Figure 5B [9] are among the most popular ones.

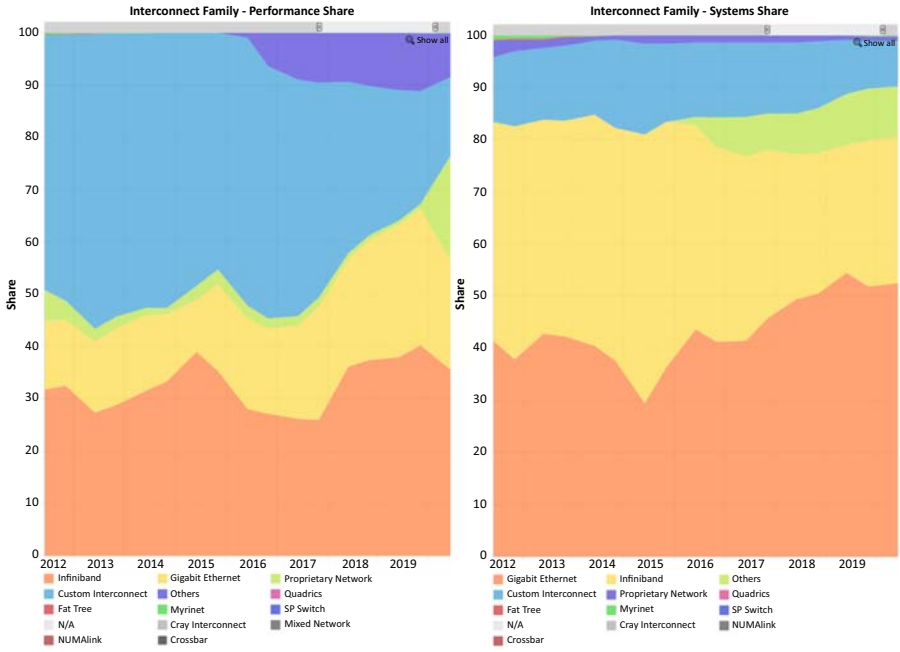

Figure 1.3 Share of interconnect families in terms of number of systems and based on the performance share.

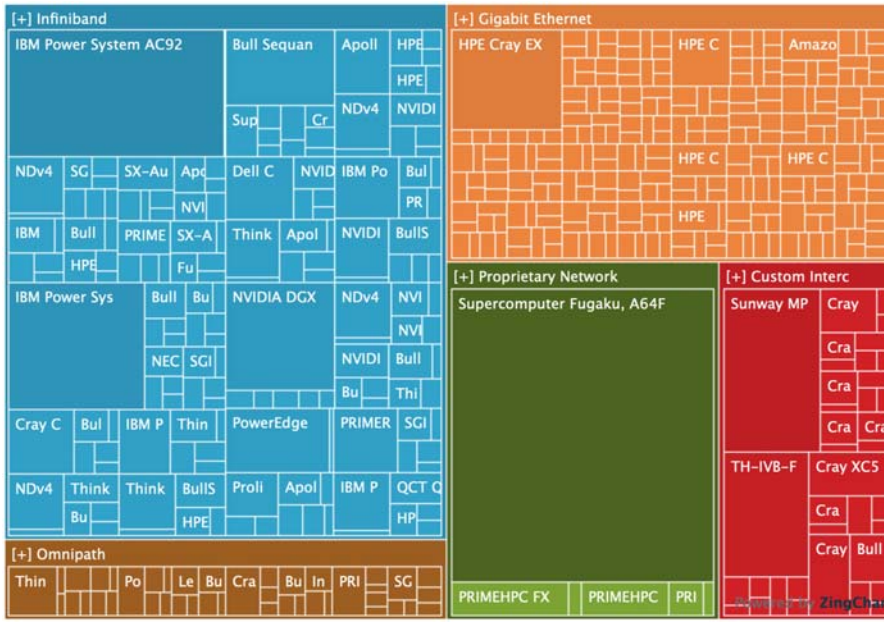

Figure 1.4 Performance share of different interconnect technologies. Many more systems use Ethernet, but they are mostly the smaller systems in the list. InfiniBand dominates the medium scale systems and proprietary + custom interconnects.

The FT topology is widely used in EDCs as it allows robust path diversity and is ideal for a multitenant DCs as it allows for a wide permutation of subnetworks with full bisection bandwidth. Combining these two aspects also allows one to, in most situations, avoid deadlocks without impeding other traffic, another feature of an FT topology. These salient features are balanced by the scalability expense of FT as one adds more levels. This is because not only that more cables are needed but they also must be longer cables, easily beyond the 50 m length mentioned before. For links of this length, often a price and technology premium must be paid. Furthermore, FT is by nature a static topology that relies on the protocol to enable path diversity, robustness, etc. Therefore scaling this topology can incur heavy latency penalties that often exclude it for certain enterprise workloads, to be discussed below. And lastly, scaling this topology requires more switches that often must contain more ports or higher bandwidth ports, where again a price and technology premium is incurred.

The last few years has demonstrated that AI and ML workloads are a necessity for the EDC market. One of the consequences of this trend is the wider adoption of traditional HPC methodologies and technologies (often this trend is referred to as hyperconvergence. However, that only refers to the hardware design and sometimes ignores the software methodologies required to fully use that hyperconverged hardware chassis). The DF topology is one such adoption as it favors not only short link reaches but also locality of data. This is because the network topology is built

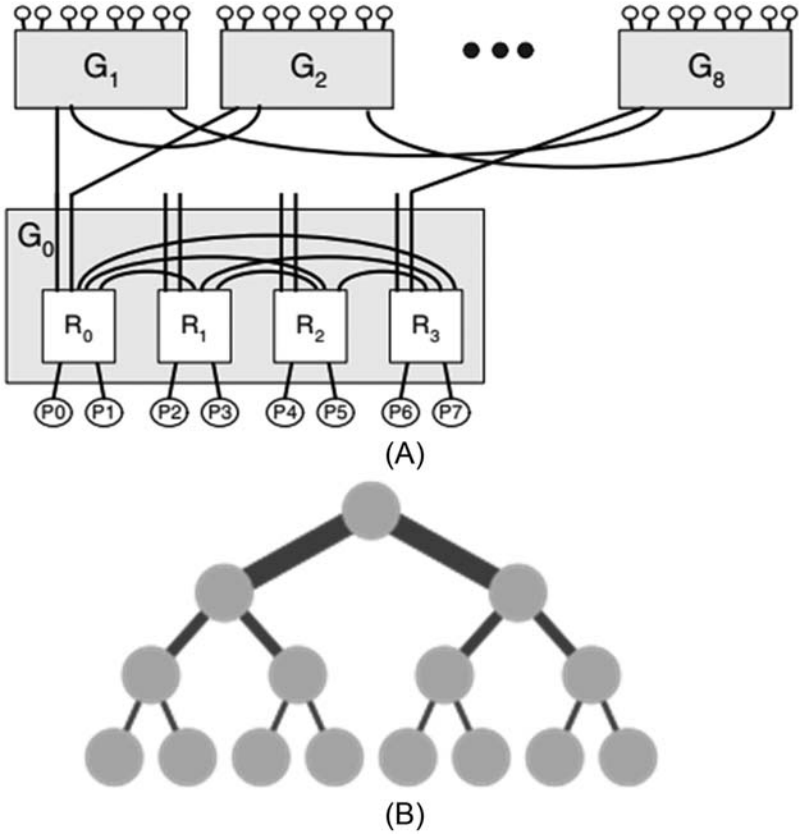

Figure 1.5 (A) Illustrative example of a Dragonfly topology, showing groups, nodes within each group, and their interconnections. (B) Illustrative example of a Fat Tree topology with three tiers of hierarchy.
Source: (A) From J. Kim, et al. "Technology-driven, highly-scalable dragonfly topology," in: International Symposium on Computer Architecture, 2008, https://static.googleusercontent.com/media/research.google.com/en//pubs/archive/34926.pdf. (B) From https://en.wikipedia.org/wiki/Fat_tree.

around the idea of clusters of compute and memory resources that are tightly coupled—both in the physical sense and in the relation of the data that the compute operate on and hold. It allows for customization of the local subnetwork while minimizing links between subnetworks. Therefore DF does not offer the degree of robustness or path diversity of FT. However, it offers a significantly better scaling option in respect of cabling price, associated hardware, and the sublinear increase in latency between subnetworks. The DF topology has been the launchpad of many other extremely compelling topologies such as butterfly, flat butterfly, and hyperx which are heavily used in the HPC community.

The EDC workloads vary widely and one criterion that can easily stratify them is their time to action. For example, a global bank with its wholly owned DC must accurately and redundantly store banking transactions and ensure absolute security of both the data and the building itself. As a result, the features of an FT topology are attractive here as the DC must work all the time, despite hardware failures or power outages. Furthermore, fixing a dead link or server is often a costly exercise due to the security protocols, so the DC must operate in a suboptimal hardware state for long periods of time. Since these transactions have a very long time to action—often they are stored and not retrieved for many years—the inefficiencies in accessing them is not a prohibitive business concern. However, the error rate with this access must be precisely zero as it often carries a financial loss or penalty.

In contrast, an advertising or marketing company who builds an augmented reality shopping app is analyzing all the data coming into its database and must be able to react in real time. Therefore latency in the time to solution is of paramount interest, so this customer might use colocated DCs (or even micro-DCs) near the customer base (e.g., a major shopping mall) to analyze that data so it can better target information sent to each user. Once in the DC, that data must be quickly inferred upon by a local neural net and often, such a neural net model is deployed on a small cluster of CPUs and GPUs, akin to the topology discussed in the example of Dragonfly. As the time to action for this example is exceptionally short, error rates can be higher: if the user occasionally receives an irrelevant ad or coupon, it is simply ignored.

Another challenge in the careful design of an EDC is the allocation of bandwidth at various hierarchy or connection points within a given topology. Often, this bandwidth is tapered in the hierarchy, referring to the decrease in per-link bandwidth as one progresses toward the server nodes. At the time of this writing, the state of the-art is a 50 Tb per second aggregate bandwidth switch realized by 64 ports at 400 G per port, with the 800 G ports in definition. This port bandwidth is the tapered down to the server level, where typical port bandwidths of 25, 50, and 100 Gb per second are used. As server bandwidths have doubled every 30 months, server bandwidths have not; this is another consequence of the slowing trend of Moore's Law. Therefore proper network bandwidth tapering will be more crucial in the years to come.

These links are realized in a pluggable form factor, the details of which are discussed further in this chapter. Such form factors enable flexibility in their aggregate bandwidth, maximum power handling, and footprint on the server faceplate. Depending on the medium, the link can reach 2−10 km using long-haul silicon photonics and 0.5−2 km using standard silicon photonics. The 10 m- to 0.5 km length is currently in flux as legacy VCSEL-based links can enable these lengths but at lower aggregate bandwidths and the industry is seeing an influx of price-competitive silicon photonics solutions that enable higher bandwidth for the same distance. For links below 10 m, passive and active copper electrical links continue

to dominate. Many of these form factor and media types adhere to various standards for the link that dictates the power consumption, transmitted and received power, bit error ratio (BER), and link reach.

1.2.3 Hyperscale data centers—cloud computing

Microsoft operates more than 200 DCs across 34 countries. The DCs are physically organized as roughly 30 megawatt installations. A typical DC might be organized with roughly 20 servers in a rack, 20 racks in a row, and around 200 rows in the full DC.

Fig. 1.6 shows a collection of servers, around 20, in a rack connected to a top-of−rack (ToR) switch with electrical cables called DACs (direct-attach copper cables). In this particular configuration, the ToR is a single point of failure for the rack and a ToR failure will render the entire rack unusable. True capital costs of the ToR and DACs, operational costs of power, and maintenance factor against availability rates for redundancy would eliminate single points of failure.

To build out a complete DC fabric, multiple racks may be networked with the two ToRs connected to Tier-1 switches (Fig. 1.7). The AOC connections are complete fiber links with transceivers mounted to the fiber during manufacturing. Deployment at the DC only requires plugging the AOC transceiver into the switch ports, eliminating any fiber mating procedures at the DC. While this can improve installation time and eliminate issues such as dust in the fiber connections, the AOCs are expensive and still prone to human failures.

The Tier-1 switches in the rows are connected to Tier-2 switches using 400 G Dr4 transceivers that use four fiber pairs per transceiver. This is a three-layer folded Clos topology that is optimized for cost and reliability while not sacrificing performance. Individual DCs are connected via longer reach optics (400G-ZR) via a regional network gateway which also provides access to the long-haul network.

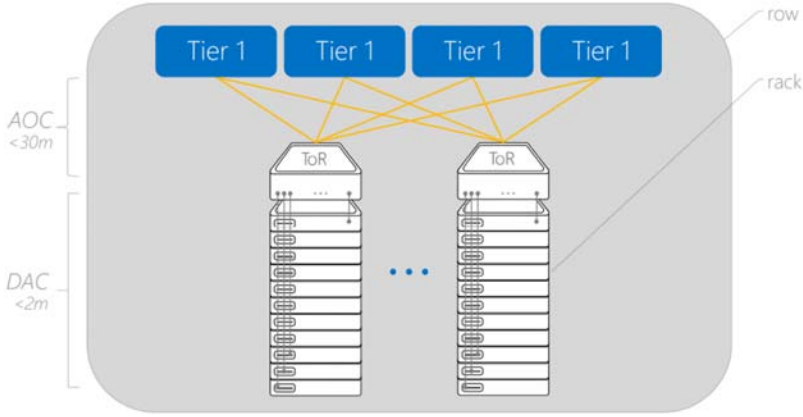

Figure 1.6 Server/ToR and Tier-1 data center switch configuration. *ToR*, Top of rack.

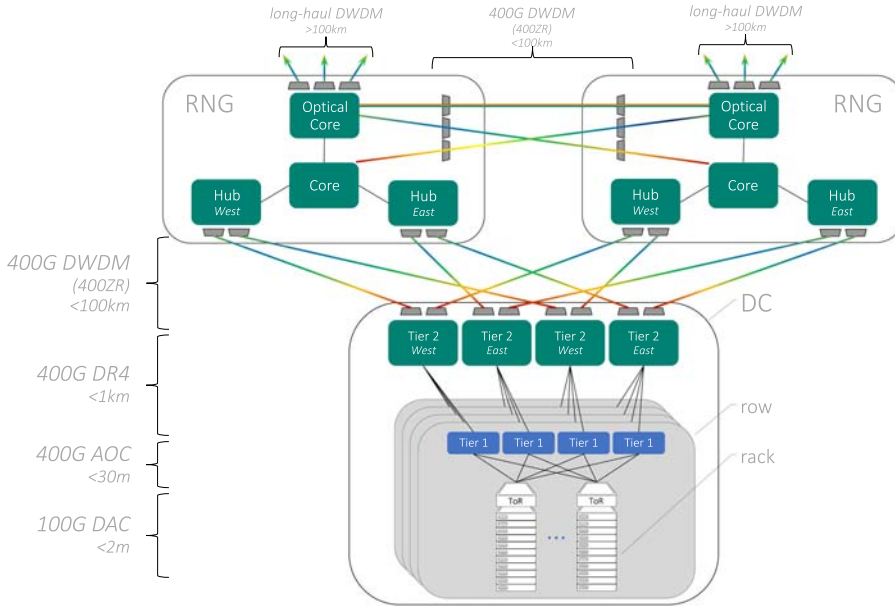

Figure 1.7 DC showing hierarchy of interconnects used and connectivity of multiple DCs via (RNG). *DC*, Data center; RNG, regional network gateway.

1.2.4 Hyperscale data centers—Web 2.0

1.2.4.1 Introduction

Meta platforms currently operate 17 global DC sites to support their suite of social media platforms. Although there are differences due to site geography, the intra-DC network follows a common design template as shown in Fig. 1.8.

Each DC site is divided up into multiple logical regions, made up of up to six data halls per region. Within each data hall, three switch tiers connect the servers: ToR, fabric, and spine, in a fat tree configuration. The network supported by these three tiers is known within Meta as the F16 Fabric [10], and is shown logically in Fig. 1.9. Server racks are arranged in multiple "pods" of 48 racks each. Server downlinks from the ToR switch are typically DAC cables for cost and power considerations. ToR uplinks from each rack connect to 16 unique fabric planes. In turn, every fabric switch connects up to 64 spine switches, interconnecting all the racks within the data hall. As each switch within a given network tier is connected with a high degree of redundancy to multiple switches in adjacent layers, F16 results in a highly fault-tolerant and scalable fabric design.

Multiple regions within the DC are connected via an aggregation layer known as "HGRID" [10], as shown in Fig. 1.10. HGRID provides the connectivity between the spine switches within the different data halls, as well as to other DC regions, and the Internet backbone.

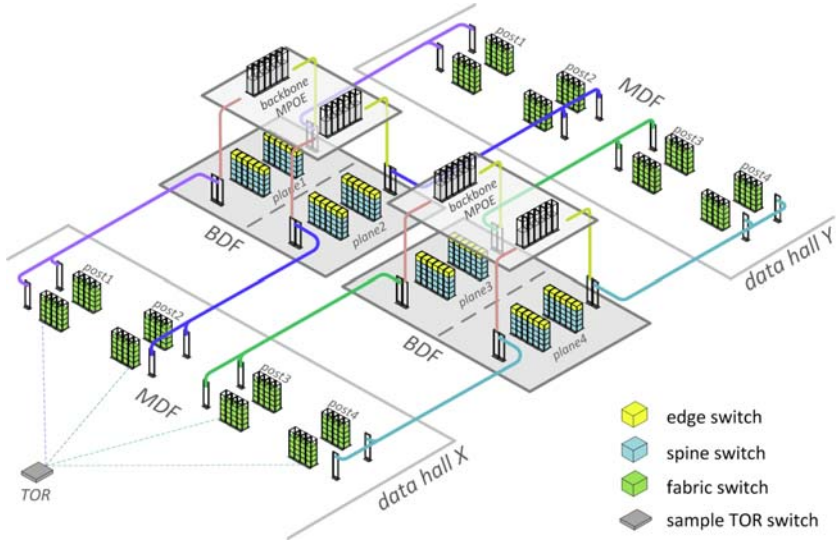

Figure 1.8 Example of Meta DC physical layout—two data halls are shown.

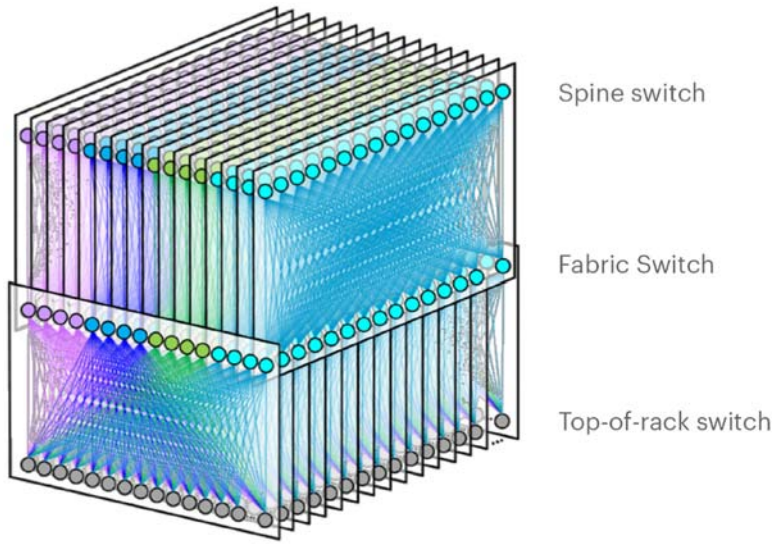

Figure 1.9 F16 fabric logical representation.

1.2.4.2 Network hardware

The F16 fabric is supported by the modular Meta designed, Minipack switch, which was contributed to OCP as an open hardware design (Fig. 1.11). This is used in locations above the ToR (fabric, spine, and HGRID). Minipack is a 12.8 Tb/s

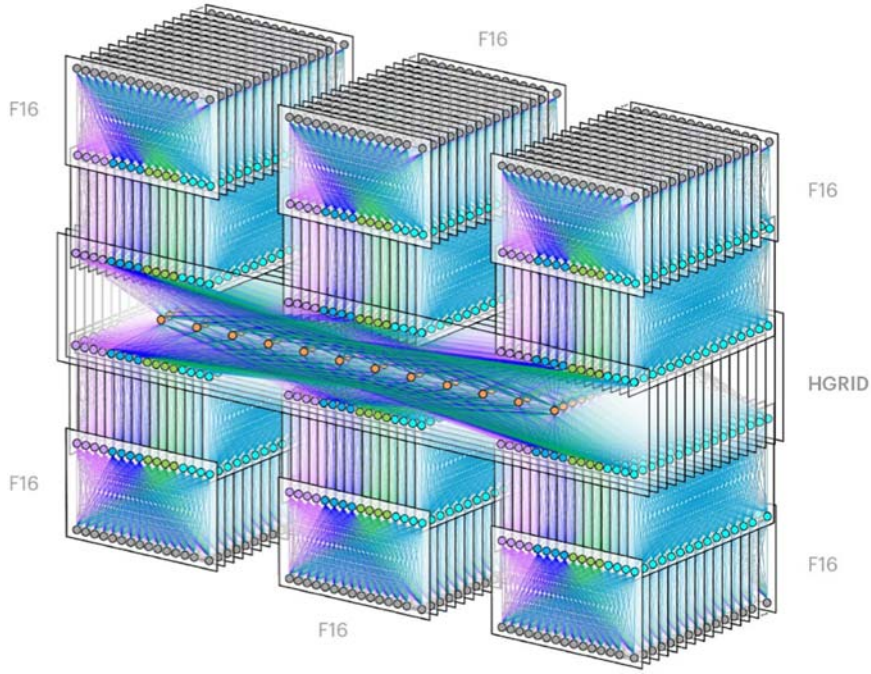

Figure 1.10 Data center region showing six data halls connected with HGRID.

Figure 1.11 Minipack 12.8 Tb/s switch, showing 128 × QSFP optical transceiver cages arranged on eight interface modules.

modular design, supporting 128 × 100 GbE ports. Use of such a radix 128 switch design is critical to enabling the F16 fabric topology. Eight vertical port interface modules can be interchanged to support a variety of optical module form factors including QSFP, QSFP-DD, and mid-board optical modules in the Meta-specified MiniPhoton [11] form factor.

1.2.4.3 Fiber and optical interfaces

Due to the physical distances associated with the DC region, maximum intrabuilding link lengths are approximately 3 km, and single-mode optical fiber is used throughout the fabric, in concert with wavelength division multiplexing to minimize the amount of total fiber required. Table 1.1 shows the generations of switch and optics hardware for use within the Meta DCs. Subsequent generations of equipment double bandwidth of switch port speed, optical module, and optical lane rates, while importantly still maintaining compatibility with the installed fiber plant and fixed infrastructure as well as enabling backward compatibility with legacy hardware in a "downspeed" mode of the new equipment. It should be noted that in the 102.4 Tb/s switch generation, power efficiency considerations suggest migration to copackaged optics (CPO) to reduce the power consumed in the switch to optics interface (see Chapter 7). Longer distance connections between DCs and to the Internet backbone require longer reach interfaces such as coherent technologies which is described in Chapter 6 on coherent photonics integrated circuits (PICs).

1.2.4.4 Workloads

In today's DCs, the vast majority of traffic is "machine to machine" in nature, rather than "machine to user" which represents only a minor fraction of the overall traffic. For the machine-to-machine traffic, these can be classified into remote storage access and remote computations. Today the former generates the most volume, while the latter causes the network the most stress. The computational traffic category typically follows a scatter-gather pattern, where a node requests an operation from a group of remote nodes. An example here would be the Ads or Newsfeed ranking head nodes that query multiple backend nodes and expect all of them to respond in bounded time in order to return the appropriate content to the user. Another prominent compute traffic example is asynchronous distributed ML training, which depends on minimizing the overall resulting communication completion time to keep the model accuracy bounded. In contrast, storage access does not generate "incast" traffic flows (which lead to congestion caused by contention) but can easily overwhelm the network with raw traffic volume, given that the ToRs are typically oversubscribed [12].

The physical network latency does not play a significant role in most cases in DC communications. The communication tail latency of the "scatter-gather" operators matters to the computational services as all pending communications need to be successfully completed before the compute operation can finish. TCP [13] represents the predominant transport protocol, when paired with the standard Linux kernel implementation. It demonstrates acceptable performance, with a few exceptions for handling the incast scenarios listed above. The latter requires specialized tuning and deployment of network-assisted signaling such as ECN [14]. As a result, DC key performance metrics are associated with internode bandwidth, and controlled tail latency under load.

Table 1.1 Meta switch and optics deployment roadmap.

Meta initial deployment year/status	Port speed (Gb/s)	Electrical lane speed (Gb/s)	Switch silicon bandwidth (Tb/s)	Switch system bandwidth (Tb/s)	Optical PMD specification	Optical lane speed (Gb/s/λ)	Optical module type		
							front-panel pluggable optics	On-board optics	Copackaged optics
2018	100	50	12.8	12.8	100G-CWDM4 (OCP)	25	QSFP-28	Mini-Photon	–
2021	200	50	25.6	25.6	200G-FR4 (OCP)	50	QSFP-56	Next-Gen OBO	–
Planning	400	100	51.2	51.2	400G-FR4 (OCP)	100	OSFP	–	CPO Gen 1
Exploration	800	200	102.4	102.4	800G-FR4	200	TBD	TBD	TBD

1.2.4.5 Machine learning hardware

As AI and ML workloads have increased as a proportion of overall DC jobs, specialized hardware is required to deliver high efficiency. Fig. 1.12 shows Meta's Zion ML platform. Zion consists of a multisocket server connected to 8 OAM (open accelerator module, [15] accelerators via an array of PCIe switches). Nodes have access to two fabrics. A coherent fabric used to interconnect the CPUs and a separate fabric for the OAMs, as show in Fig. 1.13. Connection of the Zion systems to the greater DC network is achieved via 8×100 GbE network interface controllers (NICs), via a ToR switch. The design of Zion intentionally decouples the memory, compute, and network and is optimized for memory-intensive workloads such as convolutional neural network (CNN) or Sparse neural network (NN) [16]. As

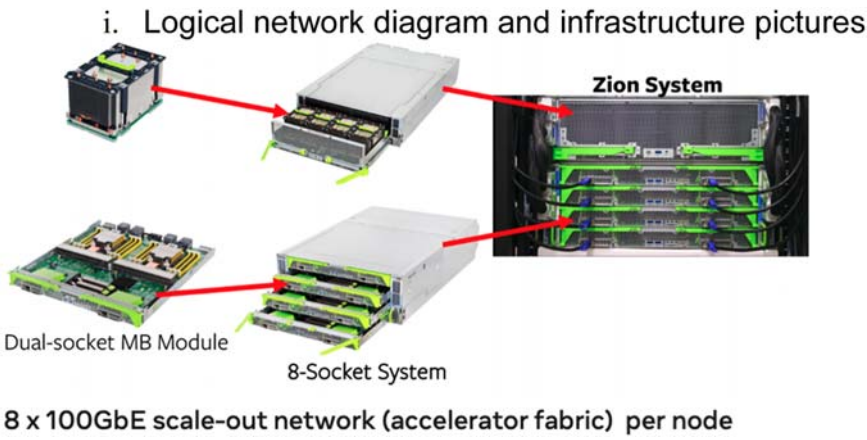
i. Logical network diagram and infrastructure pictures
Zion System
Dual-socket MB Module
8-Socket System
8 x 100GbE scale-out network (accelerator fabric) per node

Figure 1.12 Meta Zion machine learning platform.

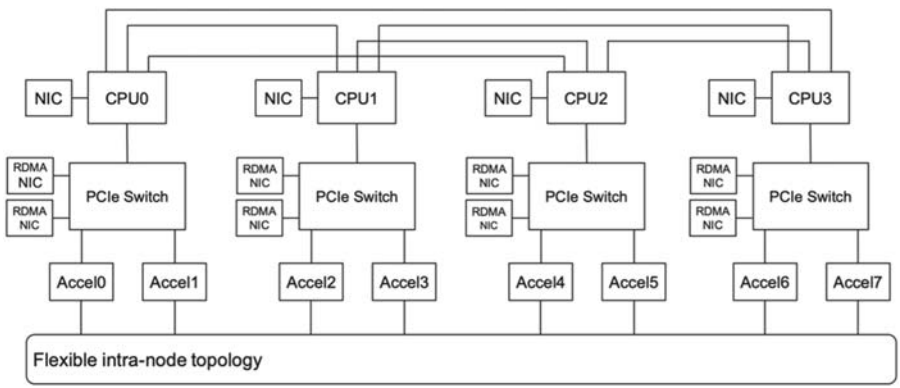

Figure 1.13 Zion connectivity and interconnect fabrics. Note this diagram shows a four-socket system as opposed to the prior generation of eight-socket design as shown in Fig. 1.11.

well as ingestion of training datasets, network connectivity is used to enable scale-out, connecting multiple Zion nodes to work on a shared job.

1.3 Optical module form factors

In today's systems, datacom interconnects primarily use front-panel pluggable (FPP) modules such as the SFP (single electrical lane), QSFP (four electrical lane), and QSFP-DD or OSFP (eight electrical lane) families [17]. A typical pluggable QSFP module is illustrated in Fig. 1.14. These modules have enjoyed wide commercial success due to their flexibility to support a wide variety of interconnect media in the same form factor; twin-ax copper cables, or single and multimode optical fiber−based interfaces supporting a variety of reaches, and are available in this form factor from a number of vendors. This flexibility enables users to employ an appropriate interconnect solution based on the application requirements of bandwidth and physical reach without having to change the installed FPP socket. In addition, pluggability enables simple field replacement in case of damage or malfunction. As a result, FPPs are the most prevalent connection on server NIC connections as well as ToR and fabric switches. Multiple generations of FPP technologies have been developed with different lane rate rates from 1 Gb/s NRZ to 100 Gb/s PAM4 signaling. In most cases, legacy compatibility within the same module family is supported, meaning users can plug a slower speed module into a higher lane speed host with appropriate rate adjustment of the host electrical interface to match the required module.

Table 1.2 summarizes commonly used optical module form factors. Selection of an appropriate optical module is dependent on the total bandwidth requirement per module (which is dictated by the desired optical interface to be supported) as well as compatibility of the host electrical interface with the desired endpoint (typically a NIC or switch). Industry standards govern the electrical, mechanical, and optical interface definitions to enable broad interoperability between module types.

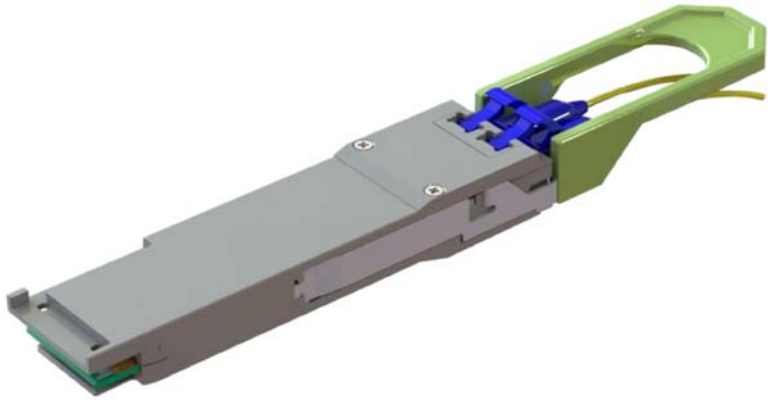

Figure 1.14 QSFP optical transceiver.

Table 1.2 Front-panel pluggable optical module form factor summary.

Module	Type	Host interface type	Maximum number of modules per RU	Electrical host interface rate (Gb/s)	Number of electrical lanes	Total module bandwidth (Gb/s)	Bandwidth per RU (Tb/s)	Width (mm)	Beachfront density (Gb/s/mm)
SFP +	FPP	Nonretimed	48	10.3125	1	10.3125	0.5	13.4	0.77
SFP-28	FPP	Retimed	48	25.78125	1	25.78125	1.2	13.4	1.92
SFP-56	FPP	Retimed	48	53.125	1	53.125	2.6	13.4	3.96
SFP-112	FPP	Retimed	48	106.25	1	106.25	5.1	13.4	7.93
QSFP +	FPP	Nonretimed	36	10.3125	4	41.25	1.5	18.4	2.25
QSFP-28	FPP	Retimed	36	25.78125	4	103.125	3.7	18.4	5.62
Mini-Photon	Mid-Board	Retimed	N/A	25.78125	16	412.5	N/A	42	9.821
QSFP-56	FPP	Retimed	36	53.125	4	212.5	7.7	18.4	11.58
QSFP-DD-56	FPP	Retimed	36	53.125	8	425	15.3	18.4	23.16
OSFP (56)	FPP	Retimed	32	53.125	8	425	13.6	22.9	18.53
QSFP-112	FPP	Retimed	36	106.25	4	425	15.3	18.4	23.16
QSFP-DD-112	FPP	Retimed	36	106.25	8	850	30.6	18.4	46.32
OSFP (112)	FPP	Retimed	32	106.25	8	850	27.2	22.9	37.07
COBO	Mid-Board	Retimed	N/A	53.125	8	425	N/A	20.0	21.25
COBO	Mid-Board	Retimed	N/A	53.125	16	850	N/A	36.0	23.61
OIF-CPO (draft)	Copackaged	Retimed	N/A	106.25	32	3400	N/A	22.5	151.11

In addition to FPP modules, products have been developed to introduce higher system density. On-board optical modules, such as those developed by the Consortium for On-Board Optics (COBO) [18], the Broadcom POD family, and also the II/VI BOA [19], are mounted directly onto the host circuit board. Although this eliminates front-panel field serviceability a mid-board approach offers placement flexibility relative to the host ASIC and in some cases simpler mechanicals which can enable higher density systems as the optical fiber can be connected via a compact front-panel fiber connector bulkhead, rather than requiring space for entire module. Use of multifiber connectors such as the MPO family in this manner effectively eliminates the system front panel as a bandwidth bottleneck due to the high fiber density which can be readily achieved.

CPO modules offer a step further in terms of density, for which integrated photonics are essential due to the size constraints of the module form factor. For prior generations of module (FPP or OBO) solutions existed consisting of either discrete solutions where individual lasers, receivers, and optical multiplexer are packaged together to form the transmit or receive optical subassemblies, or silicon PICs with either integrated, or copackaged laser engines [20,21]. In these modules, the value of a PIC-based solution compared to discrete integration is primarily economic, with the intent that a wafer scale approach would reduce cost and improve yield. The overall IO density was defined not by the optics, but the module form factor and in particular the module electrical connector. In contrast, the high density of CPO modules makes integrated photonics mandatory in order to deliver the compact form factor required to in turn enable the desired high system density, reduce the host channel length, and match the beachfront density of ASICs, which in turn enables minimization of the electrical channel between the ASIC and optical modules which permits use of lower power classes of short-reach serdes such as XSR [22].

1.3.1 Transceiver module architecture

A simplified subsystem block diagram of the transceiver types listed in Table 1.2 is shown in Fig. 1.15. Two classes of module are shown, retimed and nonretimed denoting whether the module contains a retimer device.

Taking an Ethernet-centric perspective and describing the transmit direction first, the data from the host ASIC is formatted and encoded via the appropriate IEEE 802.3-specified MAC, PCS, and forward error correction (FEC) sublayers. Electrical serdes within the ASIC then passes data over the host PCB interface and through the module electrical connector. This process is similar in nature independent of the number of interface physical lanes—the logic for multilane interfaces includes consideration for deskew and lane mapping.

For lane rates above 10 Gb/s, standard-based modules (such as the QSFP and OSFP families) have required inclusion of retimer ICs. The module retimers contain host and line-side (optical) facing serdes and optionally additional link layer logic such as PCS termination and regeneration. Where appropriate they can also support PCS format conversion including gearboxing to greater or fewer physical lanes from host to line-side. Use of retimers has proven an effective approach to manage

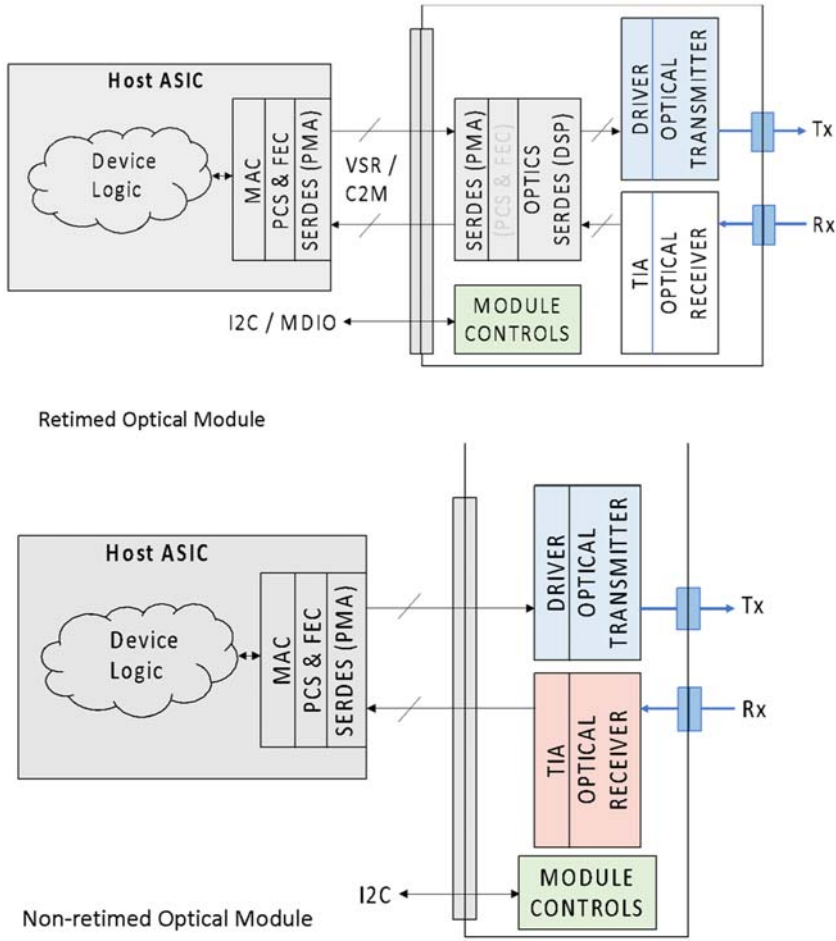

Figure 1.15 Optical module block diagrams.

signal integrity at high data rates and maintain low link error rate, as well as preserve flexibility and enable robust interoperability. However, for the latest generation optical modules, the retimer device accounts for approximately half of the overall module power consumption, which can lead to power and cooling challenges for the system designer. Recently, there has been renewed interest in nonretimed interfaces with a standards development project initiated at the Optical Internetworking Forum (OIF) for a 100 Gb/s linear interface to support nonretimed optical modules, the motivation being to support lower power and higher density interconnects [23].

After passing through the retimer (or reception from the host in the case of a nonretimed module), transmit data is then passed to the laser or modulator driver and the resulting optical signal is connected to the transmit fiber via an optical connector.

The receive direction functions in a similar manner: Received line-side data is amplified by a transimpedance amplifier (TIA) and then recovered by the line-side serdes present either in the module or host ASIC as shown in the respective figure.

Optical modules also include a control and management subsystem microprocessor to communicate with the host (typically over an I2C two wire interface for control, monitor, and alarm purposes) as well as manage the optical and electrical functions and control loops within the modules.

1.4 Interconnect figures of merit

Fig. 1.16 summarizes the interconnect technology options in terms of efficiency versus reach available to system designers today. Data was collated various silicon process nodes and lane speeds to help identify trends in terms of total interconnect efficiency, including the serdes. As expected, in terms of efficiency, short-reach electrical interfaces (either in-package, PCB or over twin-ax copper cable) offer the highest efficiencies up to distances of several meters, and system designers are able to trade off support for higher bandwidths by either addition of more lanes or use of higher rate serdes technologies.

Optical interconnects offer flexibility in terms of greater reach (up to many kilometers) but at lower overall efficiency. This is somewhat self-evident. Even for nonretimed optical modules (as shown in Fig. 1.15), a host serdes is required

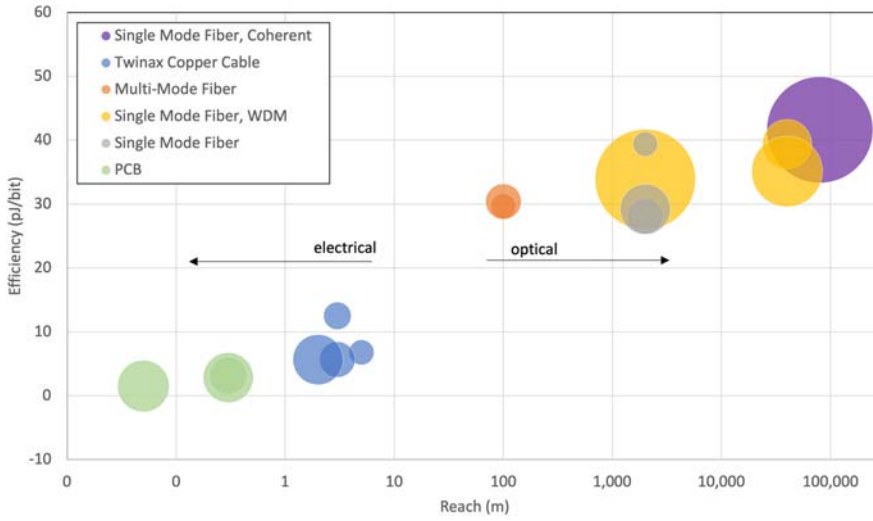

Figure 1.16 Interconnect efficiency versus reach. Size of points is relative to interconnect bandwidth per "core."

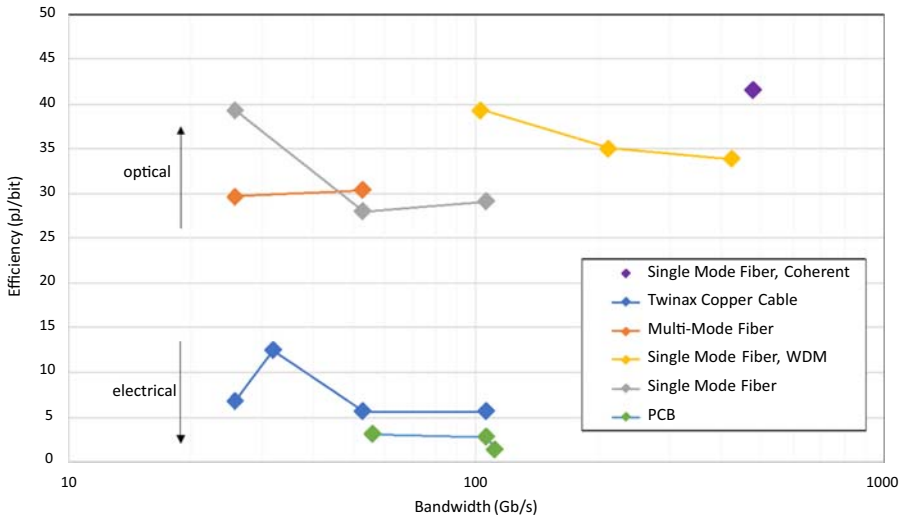

Figure 1.17 Interconnect efficiency versus bandwidth (currently available deployed technologies).

similar to an electrical interconnect; however, in the optical interconnect case, the laser, modulator drivers, TIAs, and controls consume additional power, adding to the overall module power envelope.

For electrical interfaces, as shown in Fig. 1.17, there is minimal efficiency gain in increasing lane speed by use of higher rate serdes. This trend is due to greater attenuation, reflection, and cross-talk experienced at higher frequencies, which suggest the use of a more advanced digital signal processing (DSP)-based serdes for higher baud rates 100 Gb/s and above, with more powerful equalization to offset the higher channel complexity.

There also does not appear to be a strong correlation between the efficiency and bandwidth of the optical interconnects. Multimode and parallel single-mode technologies appear comparable, followed by WDM interfaces (which are essentially multiple instances of the single-mode interfaces), and finally coherent.

1.5 Major inflection points and challenges

In this section, we consider the upcoming inflection points and challenges for data communication. The first set of challenges involves power, thermal, and the associated device and system packaging—which creates challenges with overall density. Secondly, we discuss the physical limits associated with the optical solutions in terms of optical loss and dispersion, and finally we examine resource disaggregation.

1.5.1 Power and thermal

A key focus area of integrated photonics in recent years has been to develop solutions suitable to support network switches, where the higher density afforded by integration enables placement of the optics in close proximity to the switch and eliminates the need for module retimers. Switch nodes are a logical point of optimization within DCs, as they represent the confluence points of bandwidth, so optimizing efficiency of the optical interconnects associated with these switches addresses limits in terms of overall switch system power. However, even though efficiencies of both the switch silicon and the optical transceivers have improved between generations, total powers have continued to increase as shown in Fig. 1.18. Fig. 1.18A shows that for the 102 T systems, switch ASIC power dissipation is projected to near 1 kW, which may require transition from the current approach of forced air cooling to technologies which can support higher heat thermal density such as liquid cooling.

One remedy to the diminishing efficiency of switch systems is to package the optical modules in close proximity to the network switch (known as "copackaging") and has been the focus of several industry initiatives including the OIF Copackaging project [24] as well as the Meta and Microsoft Co-packaged Optics Collaboration [25]. Copackaging enables the associated electrical channel to be shortened compared to the traditional front-panel module architecture (from approximately 250 mm to $<$50 mm), which in turn permits use of lower power and higher efficiency serdes (such as XSR) to support the switch to optics interface. An example of such an assembly is shown in Fig. 1.19, where 16 optical modules are mounted along with a switch ASIC on a common package, which is $\sim$120 mm on side. The improved efficiency afforded by this approach is forecast to reduce system power by as much as 30% compared to a FPP system. However as in this configuration the optics are mounted in close proximity to the switch, the overall thermal solution needs to be improved to accommodate both the case temperature requirements of the optical modules and switch as well as the overall higher total power dissipation of the composite copackaged assembly.

1.5.2 Resource disaggregation

As the data coming in from various edge devices (cell phones, IoT devices, security cameras) grows to be more diverse and disaggregated, placement of data structures to maintain efficient workloads becomes increasingly more difficult. To meet this challenge, concepts such as resource disaggregation and uniform data access are proposed. These two concepts both have hardware and software implications that define DCs and HPC machines. To properly disaggregate a compute or memory resource, the latency to access such a resource must be largely independent of its physical location within the physical radius of a deployed DC.

Barring for the time of flight in fiber, this means that there should not be any "step function" increases in time to access this resource due to factors such as link-level retries, switch queue, and traffic reroutes. It also means that these resources must speak the same protocol, which almost necessitates an open standard physical interface, such

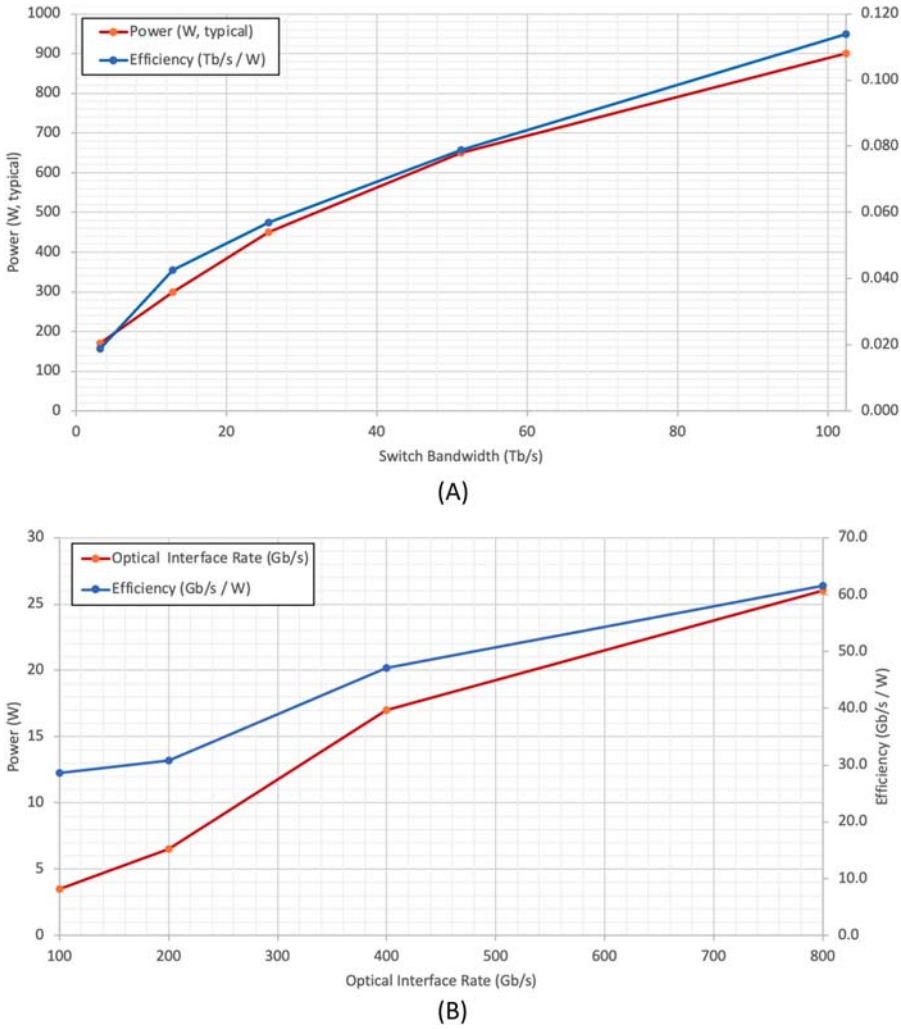

Figure 1.18 (A) Data center network switch chip power and efficiency versus bandwidth. (B) Four-wavelength of 2 km reach front-panel pluggable optical module power and efficiency versus interface rate.

as Parallel Computer Interface (PCIe Gen4, as an example). This is currently a challenge as an Intel Processor and an Nvidia GPU, both use proprietary protocols (QPI and NVLINK, respectively) which make direct intracommunication a challenge.

These capabilities necessitate that the optical interconnect have a long reach, 500 m−1 km, to span an appreciable deployment of servers and to do so with low power consumption, on the order of 10−20 pJ/bit at most. These figures of merit create attractive architectural options with the use of such interconnects. To minimize latency, these interconnects must also have high fidelity that enables a low BER of

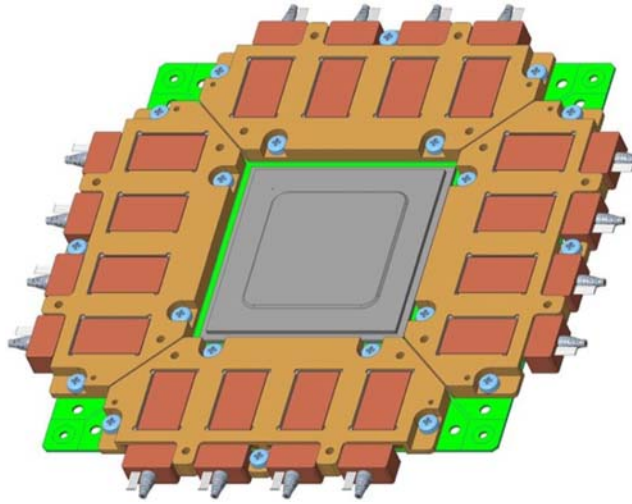

Figure 1.19 Copackaged optics and switch assembly showing 16 optical modules.

1e-9 or better, where minimal, or ideally zero, FEC can be used. This not only reduces the power but also minimizes latency that can arise from link retries and error encoding corrections. It is also of paramount importance that such interconnect technologies have a dense bandwidth per area footprint. This is needed since the wide and slow electrical links already consume a relatively large beachfront area on the package and the use of photonic interconnect that adds to this footprint would be prohibitive. As mentioned in Table 1.2, copackaged solutions enable a bandwidth per area density in excess of 150 Gb/s/mm, creating an attractive solution.

Enabling shared access to a large pool of memory (e.g.,1 PB) has certain challenges in the physical packaging of the associate media (e.g., almost 7800 128 G DIMM modules). Also, one must create a low-level protocol to enable and maintain data integrity and cache coherency, as needed. These concepts were explored and demonstrated in a project by Hewlett Packard Enterprise called The Machine [26]. This prototype machine demonstrated a flat latency access to large pools (up to 48 TB) of uniform memory that could host any types of processor, irrespective of its proprietary link, through the use of a protocol bridge. This bridge would enable a proprietary link on the client side and convert to Gen-Z [27] on the fabric side, thereby enabling the fabric to remain processor agnostic. To handle cache coherency and data integrity, Librarian, a library management software, was written that helped maintain memory pages and data security and integrity between processors and tenants.

Resource disaggregation can also be defined as the ability to deploy accelerators to a specific data location, quickly converge to a solution, and to remove that resource for redeployment elsewhere and to do it all at the software level, without physical human interactions in the DC. To realize this, the accelerator resource must have a dynamic interconnect solution that can not only define its endpoint

connections but also the bandwidth per link between it and the endpoints. In doing so, a pool of accelerator resources, such as FPGAs, can be connected in a Fat Tree topology to a cluster of CPUs and their associated data to complete a workload. Once this task is done, the same FPGAs can be redeployed to another cluster of CPUs in a Dragonfly topology, as dictated by that task and workload. Such a capability is at the forefront of advanced development today and is intended to enable the fast time to convergence that is required by future AI/ML workloads.

1.5.3 Interconnect cost targets

Fig. 1.20 shows a comparison of the approximate link costs for copper and optical interconnect typically deployed in current systems. Using the metric of cost for a normalized unit of bandwidth, or $/Gb/s, optical interconnect is approximately 11x more expensive than copper cables using retail costs. Assuming some conservative discounts for volume pricing, the relative price difference remains roughly the same.

This price comparison assumes that the copper interconnect is a DACC cable with QSFP + connectors. The optical links are assumed to be pluggable modules, also with QSFP + connectors.

Today, optical cables are used where long distances (> 10 m) span the components that need to be connected, or the bandwidth is sufficiently high that copper cables become prohibitively large. A typical DC will use copper cables within a rack to connect components to a ToR switch and optical cables to connect the ToR to other switches in the tiered network.

Ideally, optical cables could be used within the rack to improve the bandwidth between components and to support disaggregation of server nodes to allow SW

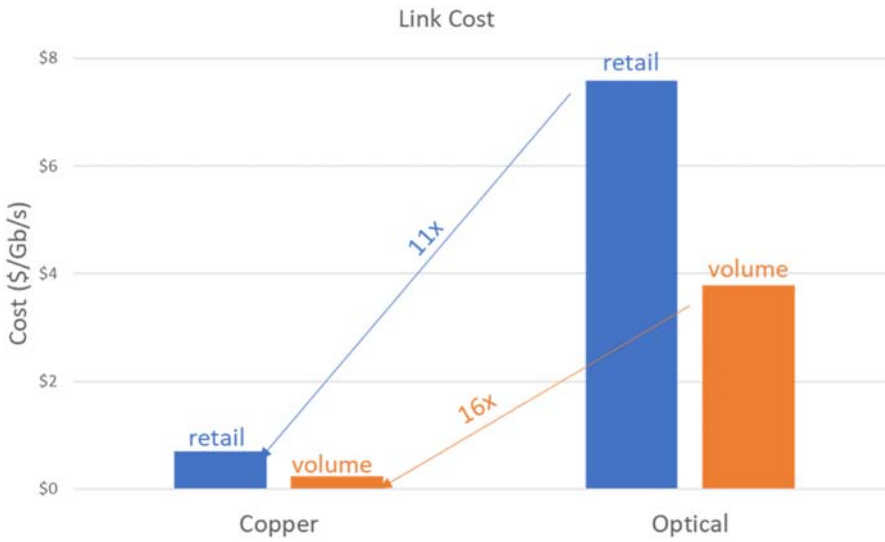

Figure 1.20 Comparison of approximate link costs for copper and optical interconnect.

composable HW configurations. One of the barriers to ubiquitous use of optical cables in DCs has been the additional cost of optical links. One goal is to develop an ecosystem and a set of technologies that will allow optical interconnects to reduce costs to match or beat the costs of copper interconnects.

CPO reduces the overhead of traditional optical pluggable modules, both reducing costs and shrinking the physical footprint of optical transceivers. Early estimates indicate that costs could be reduced by 3−5 × with CPO over pluggable modules. That would leave optics at a 2−3 × cost disadvantage relative to copper interconnect. This cost barrier could be overcome with substantially higher volumes where the economics of higher volume would drive cost down. This would come through systematic changes in optical production, including process technology flows, testing capability, and quality improvements that would be analogous to standard CMOS chip manufacturing flows.

Ultimately, the optical community should consider a cost target that would allow optics to be lower cost than copper cables. A cost target of $0.02/Gb/s, as an example, would allow optics to be the preferred interconnect for any connections that are off package (10 seconds of millimeters). This aggressive cost target would require driving the cost of every aspect of optical links to practical limits.

Next-generation network switches (51 Tb/s capacity) will require CPO to meet the physical density, power, and cost requirements. This appears to be a sufficient market-driving function to enable a CPO ecosystem to meet the demands for network switches. However, the total market of network switches is not sufficient to drive the volume necessary to achieve the cost targets that would enable optical ubiquity. The “chicken or the egg” conundrum of volume being needed to drive cost reduction to enable volume deployment is not a new one in industry and there are several performance benefits that might solve this.

Optics has several key advantages over copper interconnect, including physical reach, higher bandwidth, and smaller volumetric requirements. Strategically, if there are high-value system configurations that could not be built with copper links and fall within the capabilities of the CPO generation of optics, it might have enough volume to start to drive the cost down substantially.

Memory disaggregation could fall into this category, where the value of stranded memory in a DC is substantial enough to offset the incremental costs of optics. Importantly, many of the system configurations that are required for memory disaggregation are either complex to build with copper, or beyond the capabilities of copper because of physical reach of the disaggregation pool. Thus memory disaggregation could be the strategic lynchpin that has enough value to overcome short-term cost barriers and drive long-term cost improvements that would enable optical ubiquity.

1.5.4 Reliability

As will be covered in detail in Chapter 13, a significant reliability risk of an optical interconnect solution is the light source, often comprises III−V semiconductor materials. There are many figures of merit for lasers such as Mean Time to Failure, Failures in Time (FIT), or simply the laser lifetime as a function of operating bias.

There are also several, slightly similar, definitions of what constitutes a failure that are defined by various standards bodies. This mix of definitions leads to many non-standard implementations of where the laser should be, how it should be cooled and operated, among other concerns. For the pluggable form factor, the solution that has been adopted by the industry is to maintain an inventory on-site that enables a quick replacement in the field for a failed transceiver or AOC. This feature of an easy repair might perhaps be one of the reasons that pluggable form factors have persisted, despite the many other pain points they bring, such as very dense cable complexity.

The question of laser location of CPO has multiple facets to address insertion loss, field repairability, redundancy, and overall laser performance. Inherently, one can imagine that placing a laser closer to a high-power component immediately reduces lifetime. To be exact, it is the wide range of temperature variation seen by the laser that is the most deleterious effect. To address this, quantum dot−based gain materials are being developed that demonstrate a wider range of operation, often above 100°C. As FIT rates reduce below 100, and even 50, for some laser designs, options for including the laser inside the chassis but not necessarily on the package also become realistically possible. A growing trend in the industry is to consider the laser (or laser module) as an external power supply, akin to a voltage supply. In doing so, one can modularize the packaging for the laser supply to ease the concerns for field replacement and redundancy. At the time of this writing, standards bodies are considering an external light source standard definition to help orient the industry toward such a solution. On-package integration of lasers reduces overall link power consumption as it avoids various fiber-to-chip transitions, implying a lower operating power for a laser. This can extend lifetimes and ameliorate the laser failure issues. However, it remains to be seen that even if a laser with exceptional FIT rates in the single digits might still be integrated onto the package, repairability of such a solution is virtually impossible.

1.6 Considerations for future technology

Having described the current system applications in detail and provided context for the deployment of PICs in HPC and global DCs, the final section of this chapter identifies historical trends which can be used to guide future technology development. Based on the deployment history of multiple generations of infrastructure equipment into DCs and HPC application to date, the following considerations are put forward.

The first consideration is cost, as this plays a major role in determining the successful adoption and initial deployment of a new technology. The transition between subsequent generations of interconnect technology (typically associated with a doubling of lane speed) is often aligned with achieving parity in terms of cost per unit bandwidth. It should be noted that this is considering the total system cost holistically from end to end, including power delivery and cooling as well as the cost of changing any installed fiber infrastructure.

The next consideration which determines the successful adoption of a new technology into large DC systems is upgradeability. In many cases the new equipment represents an upgrade of an existing facility (such as a DC), and to minimize downtime and disruption to the availability of the resources being upgraded, end-users require equipment to be compatible with the fixed infrastructure which already exists. Again, this implies reuse of the power, cooling, and installed structured fiber cabling within the facility. Optical interface compatibility is also highly desirable—for example by reducing the lane rate of transceivers (or "downspeed" mode relative to full bandwidth capability), so that the next generation of lane speed WDM or parallel single-mode optics can be operated in a compatible fashion with legacy equipment. This enables upgrades to be propagated across sites with minimal disruption after initial installation, and eventual upgrade to full bandwidth can be remotely software enabled after the full-site upgrade is completed. The desire for multigeneration compatibility however stifles transformative change and places a challenge on future technologies such as new wavelength grids, or signaling formats which may be inherently incompatible with legacy technologies. For example, as dispersion and loss budgets become harder to accommodate with increasing lane speeds, a transition to low-power coherent signaling within the DC seems inevitable. However, this may require introduction in "green-field" new facilities rather than deployment within existing infrastructure due to the challenge of backward compatibility with systems in use inside DCs today.

Another consideration is compatibility or interoperability. As noted earlier in this chapter, one of the reasons for the market success of Ethernet technologies and the associated optical module interfaces has been the ability to leverage interoperable robust specifications governed by industry standards. These in turn have enabled a multivendor ecosystem which provide customers with the benefits of a strong geographically and technologically diverse supply chain which has been able to accelerate and support the high-volume requirements which the datacom industry demands. It is encouraging to note that standards projects associated with integrated photonics are already in progress, and these will be necessary for wide deployment of the new technology developments discussed throughout the following chapters.

References

[1] IEEE 802.3 Ethernet Working Group, Industry Connections Bandwidth Assessment Part II, Ad Hoc Report. <https://www.ieee802.org/3/ad_hoc/bwa2/BWA2_Report.pdf>, April 3, 2020 (accessed February, 2022).

[2] J.S. Vetter, F. Mueller. Communication characteristics of large-scale scientific applications for contemporary cluster architectures. In: Proceedings of the 16th International Parallel and Distributed Processing Symposium (IPDPS), 2002.

[3] J.S. Vetter, A. Yoo. An empirical performance evaluation of scalable scientific applications. In: Proceedings of the IEEE Conference on Supercomputing, 2002.

[4] T. Bonald, M. May, J. Bolot, Analytic evaluation of RED performance, in: Proceedings IEEE INFOCOM 2000. Conference on Computer Communications. Nineteenth Annual Joint Conference of the IEEE Computer and Communications Societies (Cat. No.00CH37064), 2000, pp. 1415−1424, vol.3, Available from: https://doi.org/10.1109/INFCOM.2000.832539.
[5] NVIDIA. NVIDIA Collective Communication Library. <https://github.com/NVIDIA/nccl>, (accessed 05.10.21).
[6] Top 500 ranked most powerful non-distributed computer systems in the world. <https://www.top500.org/> (accessed February 2022).
[7] A. Shehabi, D.A. Sarah Josephine Smith, R.E. Sartor, M. Brown, J.G. Herrlin, Koomey, N. Horner, I.L. Azevedo, and W. Lintner. United States Data Center Energy Usage Report, 2016. LBNL-1005775. <https://www.osti.gov/servlets/purl/1372902>.
[8] J. Kim, et al. Technology-driven, highly-scalable dragonfly topology, in: International Symposium on Computer Architecture, 2008, <https://static.googleusercontent.com/media/research.google.com/en//pubs/archive/34926.pdf>.
[9] Fat Tree topology Wikipedia reference, <https://en.wikipedia.org/wiki/Fat_tree> (accessed February 2022).
[10] Facebook Engineering Blogpost, Reinventing Facebook's data center network. <https://engineering.fb.com/2019/03/14/data-center-engineering/f16-minipack/>, March 14, 2019 (accessed February 2022).
[11] C.W. Berry et al., Development of on-board optics blade for hyper-scale data center fabric switches, in: 2020 IEEE Photonics Conference (IPC), 2020, pp. 1−2, Available from: https://doi.org/10.1109/IPC47351.2020.9252235.
[12] Facebook Engineering Blogpost, Introducing data center fabric, the next-generation Facebook data center network. <https://engineering.fb.com/2014/11/14/production-engineering/introducing-data-center-fabric-the-next-generation-facebook-data-center-network/>, November 14, 2014.
[13] Transmission Control Protocol Wikipedia reference. <https://en.wikipedia.org/wiki/Transmission_Control_Protocol>, February 2022.
[14] K. Ramakrishnan, Network Working Group Request for Comments, The addition of Explicit Congestion Notification (ECN) to IP. <https://datatracker.ietf.org/doc/html/rfc3168>, September 2001 (accessed February 2022).
[15] Open Compute Project OAI subgroup, "OCP accelerator module design specification_v1.5." <https://www.opencompute.org/wiki/Server/OAI>, Submitted February 23, 2022 (accessed 25.02.22).
[16] M. Anderson, First generation inference acceleration deployment at Facebook. <https://research.facebook.com/file/394856552120654/First-Generation-Inference-Accelerator-Deployment-at-Facebook.pdf>, July 8, 2021 (accessed 25.02.22).
[17] Ethernet Roadmap. <https://ethernetalliance.org/wp-content/uploads/2020/03/EthernetRoadmap-2020-Side2-FINAL.pdf> (accessed February 2022).
[18] Consortium for On-Board Optics. <https://www.onboardoptics.org/modules> (accessed February 2022).
[19] Board Mount Optical Assembly Specification. <https://ii-vi.com/product/25g-boa-board-mount-optical-assembly/> (accessed 25.02.22).
[20] S. Fathololoumi, K. Nguyen, H. Mahalingam, M. Sakib, Z. Li, C. Seibert, et al., 1.6Tbps silicon photonics integrated circuit for co-packaged optical-IO switch applications, in: Optical Fiber Communication Conference (OFC) 2020, OSA Technical Digest (Optica Publishing Group, 2020), paper T3H.1.

[21] T. Pinguet et al., High-volume manufacturing platform for silicon photonics, in Proc. IEEE, 106, 12, pp. 2281–2290, Dec. 2018, Available from: https://doi.org/10.1109/JPROC.2018.2859198.
[22] Common Electrical Interface XSR Project, Physical and Link Layer Working Group Chair Dave Stauffer. <https://www.oiforum.com/technical-work/current-work/#cei-xsr> (accessed 25.02.22).
[23] Common Electrical Interface 112G Linear Project, Physical and Link Layer Working Group Chair Dave Stauffer. <https://www.oiforum.com/technical-work/current-work/#CEI-112G%20linear> (accessed 25.02.22).
[24] Optical Internetworking Forum Co-Packaging Framework <https://www.oiforum.com/technical-work/hot-topics/co-packaging/>.
[25] Co-packaged Optics Collaboration website. <http://www.copackagedoptics.com/> (accessed February 2022).
[26] Memory Driven Computing, The Machine. <https://www.hpe.com/us/en/newsroom/content-hub/memory-driven-computing.html> (accessed February 2022).
[27] Gen-Z Consortium. <https://genzconsortium.org/> (accessed February 2022).

This page intentionally left blank

Integrated lasers for data center silicon photonic-integrated circuits

2

Chao Xiang[1,2] *and John E. Bowers*[1]
[1]Department of Electrical and Computer Engineering, University of California, Santa Barbara, CA, United States, [2]Department of Electrical and Electronic Engineering, The University of Hong Kong, Hong Kong, China

2.1 Introduction

Integrated photonics plays an important role in data center interconnects and communications. With the rapid expansion of data center scales, the number of adopted integrated photonic components increases dramatically over the past decade and the trend will continue and expand. For all photonics-related applications, lasers are key devices and semiconductor lasers are undoubtedly the solution for data center applications due to the unparalleled advantages in size, cost, and power consumption, which are critical factors driving the development of hyper-scale data centers [1].

Mature semiconductor lasers in data communications use group III−V materials including indium phosphide (InP), gallium arsenide (GaAs), and other lattice-matched quaternary alloys including InGaAsP, InGaAlAs, etc., covering telecom O-band or C-band. III−V-based photonic-integrated circuits were the mainstream technology in long-haul optical communication systems. However, silicon photonics is gaining increasing attention in data center interconnects nowadays due to its overall lower cost, lower power consumption, and higher scalability. High-yield and high-performance silicon photonic devices, including modulators and photodetectors (Si/Ge), are desired in the fast-growing hyperscale data center applications and are equipped with the additional advantages of low-cost and high-throughput CMOS manufacturing process. As a comparison, the laser sources, which are absent from monolithic Si due to its indirect bandgap, need to rely on III−V gain materials including InP or GaAs. As such, an important debate for silicon photonic applications in data centers is whether the laser needs to be discrete or integrated with Si PICs in any forms. Several approaches of laser implementation and their current status, issues, and future directions will be covered in this chapter.

In addition, we will introduce the working principles, performance metrics, and issues of several different laser approaches and devices in data center applications. We will also showcase several representative state-of-the-art laser devices and discuss the future perspectives.

Integrated Photonics for Data Communication Applications. DOI: https://doi.org/10.1016/B978-0-323-91224-2.00009-6
© 2023 Elsevier Inc. All rights reserved.

2.1.1 Types of sources needed for datacom

The choice of lasers for data communications depends on many factors, such as the overall link architecture, capacity, and reach as well as component cost and power consumption. For example, the laser configuration can be generally categorized as vertical-cavity surface-emitting lasers (VCSELs) and in-plane lasers. VCSELs dominated in short-reach (<200 m) 850-nm multimode fiber-based optical networks. As new data center builds have focused on longer reach to accommodate the scale, single-mode fiber links are preferred with the use of in-plane/edge-emitting lasers. This chapter will focus on in-plane lasers suitable for the growth of data center interconnect in both the speed and reach, while information on VCSELs can be found in other literature [2].

There are several methods to differentiate semiconductor lasers based on their properties, and one important criterion is their spectral characteristic (Fig. 2.1A). For example, multiwavelength lasers can be formed with simple cleaved facts, but to be used in higher capacity wavelength-division multiplexing (WDM) architectures, there are strict requirements for the performance of individual spectral lines. Single-wavelength lasers are often required and rely on wavelength-selective filters. Tunable-wavelength lasers are used in dynamic wavelength-routed architectures that are currently more common in telecom networks and are moving into data centers with the increase in both interconnect reach and capacity. The laser spectral characteristics depend on the mode filtering element and its position relative to the longitudinal cavity modes. Fig. 2.1A summarizes this effect and three types of the most widely used semiconductor lasers with different spectral characteristics: (1) single wavelength lasers; (2) multiwavelength lasers; and (3) tunable-wavelength lasers that will be discussed in this section. For each type of laser, there are several laser cavity designs for the desired mode filtering effect as shown in Fig. 2.1B.

2.1.1.1 Multiwavelength lasers

2.1.1.1.1 Fabry-Pérot lasers

Fabry–Pérot (FP) lasers are one of the simplest forms of lasers because they do not require a special cavity design, and the cleaved or polished facet forms a natural partially reflective mirror for laser oscillation. As the laser cleaved/polished facets show little wavelength selectivity, FP laser outputs are multiwavelength, and the wavelengths red-shift as a result of the bandgap shrinkage with increased laser gain current and heating-induced refractive index change.

Due to their simplicity and low cost, FP lasers are often used in communication systems where a direct modulation scheme is applied. In view of the multimode spectral characteristics and incompatibility with WDM systems, FP lasers normally find applications with limited reach (~100 m) and limited modulation speed (<10 GHz). The longitudinal modes in an FP laser are not phase locked together and the power in a particular line is not stable. FP lasers are not suitable for multiwavelength WDM systems where each wavelength is dedicated to a particular data stream. As a comparison, stable multiwavelength lasers (comb lasers) that are compatible with WDM systems are discussed later in this section.

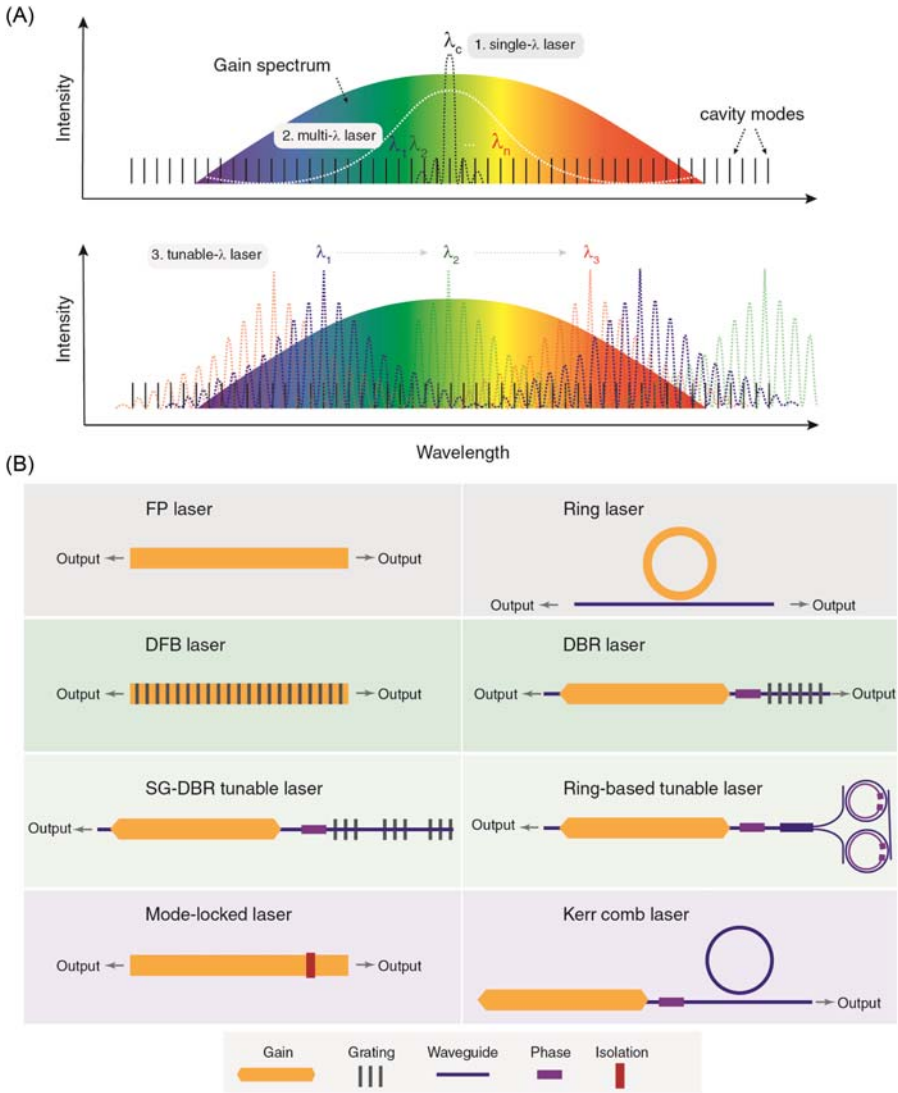

Figure 2.1 Based on the wavelength selection filter, lasers are categorized as (1) single-λ laser, (2) multi-λ laser, and (3) tunable-λ laser. The wavelength filters for each type are shown as dotted lines.

2.1.1.1.2 Ring lasers

Ring-resonator-based lasers utilizing whispering gallery mode in a ring-shaped cavity for laser oscillation offer compact solutions for low-power-consumption data communications. They have no requirements for cleaved-facet lasers or grating-based cavities. The lasing spectrum is multi-longitudinal mode with ring free-spectral range-defined mode spacing. However, small diameter lasers have a main

lasing mode with decent side mode suppression. Heterogeneously integrated ring lasers can be integrated with metal-oxide-semiconductor tuner for ultra-efficient wavelength tuning and high-speed direct modulation [3].

2.1.1.2 Single frequency lasers

2.1.1.2.1 Distributed-feedback lasers

Distributed-feedback (DFB) lasers are arguably the most widely used single-wavelength laser in optical communication systems, as the single mode operation (narrower spectral width) results in less dispersion penalty and longer transmission distance than FP lasers. They have single-wavelength output and are also suitable for WDM systems in which DFB laser arrays are employed. The working mechanism of a DFB laser is based on a Bragg grating that provides single-wavelength filtering due to the grating bandgap, and this filter element is placed along with the laser gain section. The Bragg grating period determines the lasing wavelength, so for WDM systems, DFB arrays with different grating period offsets enable different lasing wavelengths to align with standardized wavelength grids.

The single-mode operation of DFB lasers, is sensitive to the phase relations of the laser cavity. In many cases, quarter-wavelength shift section or highly-reflective/anti-reflective (HR/AR) coatings need to be included to provide nondegenerate phase relations favoring single longitudinal mode lasing. DFB lasers can be used in both direct-modulation and external-modulation configurations. Direct modulation provides a simpler, lower-cost solution, but the modulation speed is limited and requires high current to increase the bandwidth. Recent developments in external feedback control enabled over 100 Gb/s direct modulation of DFB lasers and expanded DML applications into single-lane beyond 100 Gb/s systems [4]. For external modulation in long-haul communications, high-power, high-SMSR (side mode suppression ratio), and high-coherence DFB lasers are integrated with Mach−Zehnder modulators, electro-absorption modulators or microring modulators (MRMs). Details on the modulators could be can in Chapter 3.

2.1.1.2.2 Distributed Bragg reflector lasers

Another type of grating-based single-wavelength laser is the distributed Bragg-reflector (DBR) laser. Unlike DFB lasers whose gratings are coupled with the laser gain section, the gratings of DBR lasers are separated from the laser gain, which enables potential separate tuning of the grating index and thus the lasing wavelength. In addition, as the grating is a passive device external to the cavity, it offers opportunities in laser performance engineering, such as reducing the laser optical linewidth and relative intensity noise (RIN) [5]. Due to the separation of laser gain and the feedback grating, there is common mode-hop behavior in DBR lasers. To ensure single-mode operation, a phase section between the laser gain and Bragg grating is normally included (Fig. 2.1B).

2.1.1.3 Tunable lasers

Tunable lasers are often used in communications systems that require dynamic channel allocations and reconfigurability. In cases where the cost of tunable lasers

is comparable to that of nontunable lasers, they can replace DFB lasers. Another important feature of tunable lasers is that they normally have narrow spectral linewidth, which is ideal for coherent communications and advanced high-order modulation formats. Several cavity configurations can enable wavelength tunable feedback for the laser and the wavelength tuning is mostly based on thermo-optical or electro-optical effect.

2.1.1.3.1 Sampled-grating distributed Bragg reflector lasers

Sampled-grating DBR (SGDBR) lasers are based on the Vernier effect from sampled gratings. The amplitude or the phase of the gratings in the front and rear mirrors is periodically modulated. The reflection spectra see periodic maximum amplitude response with a spacing determined by the inverse of the distance between adjacent grating sampling windows. Front mirrors and back mirrors have offset sampling length which results in the Vernier effect leading to a wide wavelength tuning range.

SGDBR lasers are normally used in InP-based widely tunable lasers where multiple regrowth to form the laser cavity structure needs to be performed. The wavelength tuning is achieved by tuning the grating index of the feedback mirrors. Fine and continuous wavelength tuning is accomplished by tuning the front-mirror and back-mirror index at the same rate. On the other hand, wavelength channel switching (wide-tuning) can be achieved by differential effective index change of the front- and back-mirrors. A typical tuning range of SGDBR lasers is 40 nm or wider and the tuning is normally based on electrical current injection [6].

2.1.1.3.2 Vernier ring lasers

Compared with grating-based cavities with fixed optical bandgap (central wavelength), ring resonators are featured with wide wavelength tunability, which are ideal to be used in tunable lasers. However, a single-ring resonator does not provide enough feedback strength difference between resonances and lasers based on single ring-resonator normally having limited SMSR. A design based on two-ring resonators solves this issue by introducing an additional ring resonator with slightly different dimensions to suppress feedback from the neighboring resonances and the resultant tuning range is thus significantly enhanced by the Vernier effect [7].

The high-quality factors of ring resonators provide another benefit of ring-resonator-based tunable lasers. As an example of on-chip external cavity, ring resonators could greatly extend the passive cavity length and dramatically decrease the laser linewidth. Si or Si_3N_4 passive ring resonators can be integrated with III−V gain materials for high-performance widely tunable lasers [8]. For such lasers, the tuning is normally based on thermo-optic resistive heaters to minimize the intracavity loss and maintain the high Q factors of the ring resonators, which are critical to the laser performance.

2.1.1.4 Comb lasers

To deploy WDM schemes in high-capacity communication systems, multiple signal channels at predefined fixed wavelength grids are required. A laser array using multiple single-wavelength lasers is an option but has the complexity and cost of controlling multiple laser wavelengths. A single laser with multiwavelength output has the

advantage of fixed channel spacing to simplify the laser wavelength control and significantly lower the cost and power consumption of solutions requiring a large number of channels. Using a single multiwavelength laser instead of multiple single-wavelength lasers also gains a broadband phase coherence between the laser lines, which enables new techniques, such as injection locking, to manipulate the laser line properties. Two types of comb lasers could be used in data communications: mode-locked comb lasers and Kerr nonlinear frequency combs.

2.1.1.4.1 Mode-locked lasers

A mode-locked laser is a coherent form of a multiwavelength laser. For FP lasers, the multiple longitudinal modes oscillate independently. In contrast, the longitudinal modes of mode-locked lasers follow constant-phase relations and are thus coherent with one another. Mode-locked lasers output short pulse trains in the time domain, which correspond to the comb line spacing in the frequency domain. For mode-locked semiconductor lasers, the mode spacing is determined by the group index of laser cavity and cavity length and normally falls in the microwave or millimeter wave range. The mode spacing can thus be compatible with standard DWDM (dense wavelength division multiplexing) grids (50 GHz) or be engineered to comply with different system configurations. In systems where over 100 GHz spacing is desired, it is challenging to rely on fundamental mode-locking. High-order harmonic mode-locking scheme can be utilized to suppress the most neighboring modes and achieve extended mode spacing.

A typical mode-locked semiconductor laser consists of a laser gain and a saturable absorber section, which can be implemented using the same III−V materials (Fig. 2.1B). The laser gain current and the applied reverse bias on the saturable absorber determine the laser operation conditions for ideal mode locking. Self-mode locking is observed in quantum dot (QD) single-section lasers without saturable absorbers due to the large four-wave mixing coefficient and inhomogeneous distribution of QDs [9]. Single-section self-mode-locked lasers could provide higher wall-plug efficiency (WPE) due to the absence of the saturable absorption.

2.1.1.4.2 Kerr comb lasers

Another type of optical frequency combs (OFCs) is based on the nonlinear Kerr effect in optical resonators. Of particular interest are the OFCs generated in microresonators, termed microcomb. Cascaded four-wave mixing at a wide bandwidth occurs when the high-power pump laser matches the resonance of a dispersion-engineered high-Q microresonator, resulting in a large number of generated laser lines with wavelengths spacing determined by the resonator dimensions. Apart from the applications in optical metrology and spectroscopy, Kerr OFCs with favorable properties including high-channel-count, high-coherence, and low-phase noise are suitable for high-capacity optical communication systems. Depending on the configuration and dispersion properties of microresonators, the generated comb could have distinct spectral behaviors, including dissipative Kerr solitons, dark pulses, photonic molecules, and so on.

Recent progress in integrating semiconductor lasers with nonlinear microresonators is one critical step toward future deployment of Kerr OFCs in data communications [10−12]. Different types of Kerr OFCs with variable resonator designs or operation principles are also explored to improve the power conversion efficiency [13]. Simple and reliable OFC generations are demonstrated with turnkey operation, which is critical in practical data communication applications [14].

2.1.2 Laser solutions

Silicon photonic-integrated circuits provide high-performance, low-cost solutions for data communications and there are different approaches to utilize III−V-based semiconductor laser sources (Fig. 2.2). The deciding factors range from device-level material compatibility and reliability, to system-level architectures and cost. Lasers can be disaggregated from the Si PICs and coupled through fiber connections. On the other hand, a laser and/or semiconductor optical amplifier (SOA) integrated with Si PICs has advantages of compactness and robustness. The overall cost of integrated lasers with Si PICs will also show advantages when the number of lanes increases further. In this section, we discuss both the use of disaggregated sources and integrated sources that have different integration approaches with Si PICs.

2.1.2.1 Disaggregated sources

Disaggregated laser sources with direct current modulation provide a low-cost solution for optical transmitters without the need for external modulators. By modulating the applied current on the laser diode, the laser output intensity is modulated as a result. The bandwidth or modulation speed of DMLs is normally below 50 GHz,

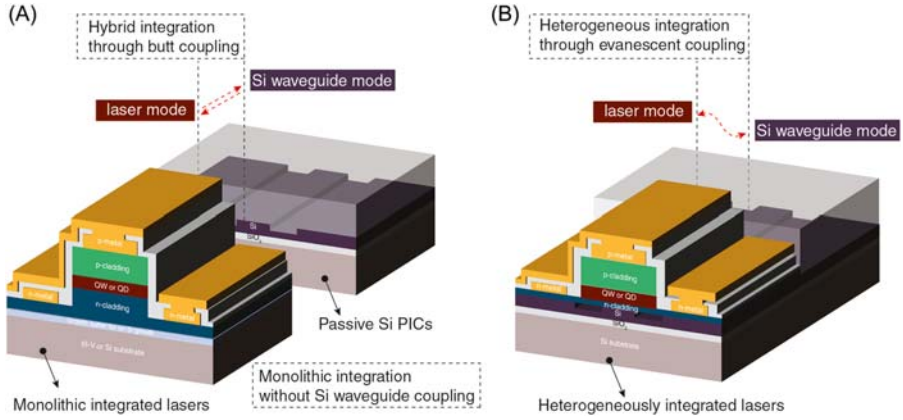

Figure 2.2 The schematic configurations of lasers with different types of integration approaches with Si PICs. The monolithic lasers need to be coupled to Si PICs through a butt-coupled configuration to form hybrid integration. Heterogeneous-integrated lasers are capable of direct on-chip light coupling to Si PICs.

although recent works have reported 65−75-GHz bandwidth for laser designs exploiting detuned loading, photon−photon resonance, and in-cavity frequency modulation−amplitude modulation conversion [4]. For 1.3-μm DMLs, InGaAlAs multiple-quantum-well gain region is widely used as it provides higher carrier confinement than InGaAsP. To reduce the laser bond pad capacitance, benzocyclobutene (BCB) layers can be used to increase the laser RC cutoff frequency [15].

Disaggregated lasers are connected to silicon PICs using optical fibers. Depending on the transceiver architecture, the disaggregated III−V lasers can be figured into different architectures. For example, in parallel single mode (PSM) architecture, a single laser output is split into several channels (four channels in PSM4) for individual signal modulation, followed by individual fibers in the transmission link. In comparison, in CWDM4 (coarse wavelength division multiplexing-4) configuration, four lasers at different wavelengths with 20-nm spacing are required and fed into a single fiber. The lasers can be either directly modulated or externally modulated.

Advanced configurations can be deployed with a mix of PSM and CWDM to increase the link capacity. Each DFB laser in a DFB array can be split into several fibers and as a result, multiple fibers will carry WDM signals in parallel. For each channel, signal modulation happens after the signal multiplexing, and MRMs provide a compelling solution for simultaneous signal filtering and modulation without the need for wavelength multiplexers/demultiplexers.

Disaggregated laser sources are also termed as “remote laser sources” or “external laser supplies.” They have to rely on fiber connections to Si PICs and thus allow decoupling of the thermal environment of the laser and Si PICs. In co-packaged optics (CPO) configurations, remote laser locations would also avoid the heat dissipation to the switch fabrics. The remote laser source configuration is determined by the system capacity, PIC split ratio, and also the laser power at the maximum module case temperatures. However, the edge density of coupling fibers to Si PICs is limited and further scaling of number of fibers is limited. In addition, the laser to Si PIC coupling loss and nonlinear absorption loss in the silicon waveguide will also limit the power budget when more lanes are required.

2.1.2.2 Integrated sources

2.1.2.2.1 Hybrid integration

Hybrid integration refers to the integration approach in which separate chips are coupled to each other in close proximity with fine alignment or attachment. The III−V gain chip can be coupled to Si PICs through direct butt-coupling, photonic wire-bonding, free space coupling, etc. (Fig. 2.3). Active alignment or precision etch stops could be used to align the gain chip to Si PICs. While the III−V gain region is still separate from the Si PIC chip, same as disaggregated sources, this approach allows the integration of III−V laser gain with Si-based laser cavities. Hybrid-integrated external cavity lasers, including DBR laser, and ring-resonator-based tunable lasers have been demonstrated with favorable performance thanks to the low loss Si external cavities.

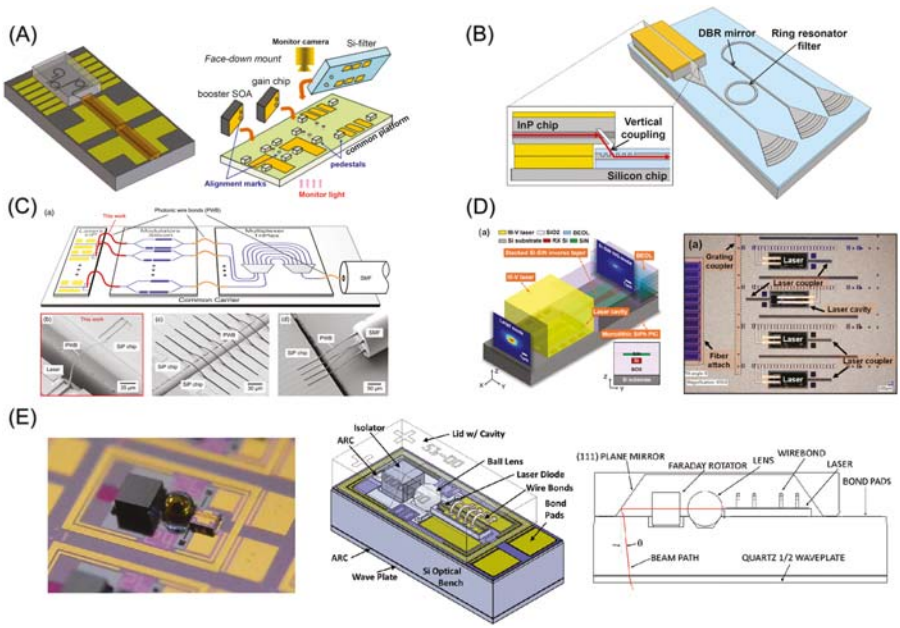

Figure 2.3 Examples of hybrid laser-Si PIC integration techniques. The hybrid integration is based on: (A) passive alignment [16]; (B) metal bonding [17]; (C) photonic wire bonding [18]; (D) passive alignment aided by SiN spot size converter [19]; and (E) Luxtera's hermetic silicon Laser Micro Package [20].

The external cavities could also be Si_3N_4 waveguides, which could enable orders of magnitude lower loss than state-of-the art Si. Narrow-linewidth lasers based on Si_3N_4 Bragg gratings and ring resonators have achieved rapid progress. In addition to the external-cavity configuration, self-injection locking (SIL) can be implemented using III−V lasers (FP or DFB) and high-Q ring resonators to dramatically reduce the laser linewidth by seven orders of magnitude [21]. This is potentially important in digital signal processing (DSP)-free coherent optical links.

Flip-chip bonded lasers are also a type of hybrid integration. Similar to butt coupling, the laser alignment with Si PICs needs to be precisely controlled to minimize the coupling loss. Spot size converters or prisms could be used in different demonstrations [16−19].

2.1.2.2.2 Heterogeneous integration

Heterogeneous integration leverages wafer bonding technology to directly integrate efficient III−V gain materials with Si PICs for lasers (Fig. 2.4). This has become the most viable technology to enable lasers directly integrated on silicon waveguides [23,24]. Heterogeneous integration process starts with Si waveguide patterning and then III−V materials are bonded onto the Si wafer, followed by III−V substrate removal and subsequent laser processing. Multiple III−V epis can be bonded on the same wafer for optimized functional devices (lasers, modulators,

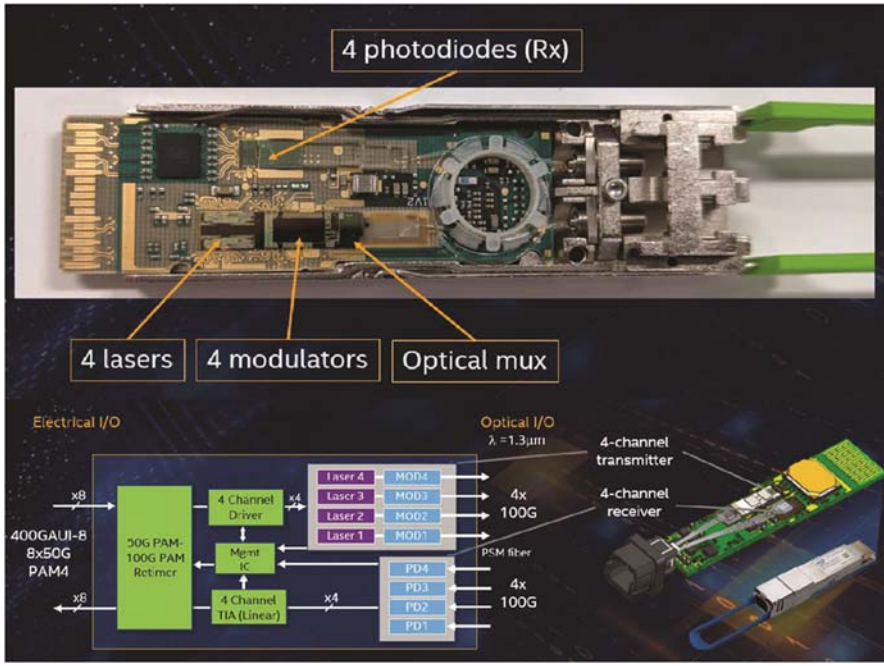

Figure 2.4 Intel's nonhermetic packaged heterogeneous InP/Si transceiver picture and configuration [22].

photodetectors, etc.). One important advantage is that laser gain epis with different bandgaps could be bonded on the same Si wafer and enable efficient laser sources for a wide variety of wavelengths. This plays an important role in architectures where laser wavelengths vary a lot across the channels, for example, CWDM8.

Several process variations can be made in heterogeneous integration. For instance, micro-transfer printing is a method to selectively "pick and place" a partially processed III−V coupon which requires accurate alignment. As a comparison, wafer-to-wafer bonding or die-to-wafer bonding of blank films has no requirements on the fine alignment at the bonding step since the final device alignment is performed at the subsequent lithography steps. Complex heterogeneous Si PICs can enable complete optical functionalities on chip. The CMOS process and capability of using 300-mm Si substrates promised the mass-scalability of heterogeneous integrated lasers and Si PICs.

Heterogeneous integration also completely eliminated the fiber coupling loss between lasers and Si PICs and could increase the link power budget. Coupling losses of under 1 dB with reflections of less than −40 dB have been demonstrated [25]. Moreover, heterogeneous integration of lasers directly on Si PICs enabled complex devices with the highest density. With the further increase in number of multiplexed channels, heterogeneous integrated lasers offer the benefits of lowest cost and complexity. The number of working channels is not limited by the chip

coupling density for fibers or layout constraints for hybrid butt coupling, in the context of spatially multiplexing the number of lanes.

In addition, heterogeneous integration allows the bonding of different III−V epi (quantum well and QD) on a common Si wafer to include the optimized epi selection for III−V/Si modulators and photodetectors, in addition to lasers and SOAs for fully functional III−V/Si PICs [26−28].

2.1.2.2.3 Monolithic integration

Monolithic integration on silicon utilizes the same Si substrate in the entire growth and processing sequence. For monolithic lasers on silicon, this eliminates the expensive GaAs or InP substrates in the process and offers the lowest cost. Directly growing lasers on 300-mm Si substrate also offers the best scalability. However, due to the antiphase domains, lattice constant mismatch, and thermal expansion mismatch between III−V materials and Si, epitaxial III−V lasers on Si have been challenging. QW lasers grown on Si substrates have had poor performance (Fig. 2.5A) and extremely poor lifetimes, just 1000 hours at room temperature.

Epitaxial growth of quantum dot materials on silicon has been discovered as a promising approach and one main reason is that QDs are much less sensitive to dislocations on silicon. With developments in epi design (buffer layers, defect blocking layers), growth development (thermal cycling), geometry engineering (high confinement in SCH), and so on, the QD laser lifetimes have improved significantly, and so does the rest of QD-on-Si laser performance (threshold, slope efficiency, power, maximum temperature) [30]. The next research direction is to effectively integrate the QD laser output with Si PICs. State-of-the art results will

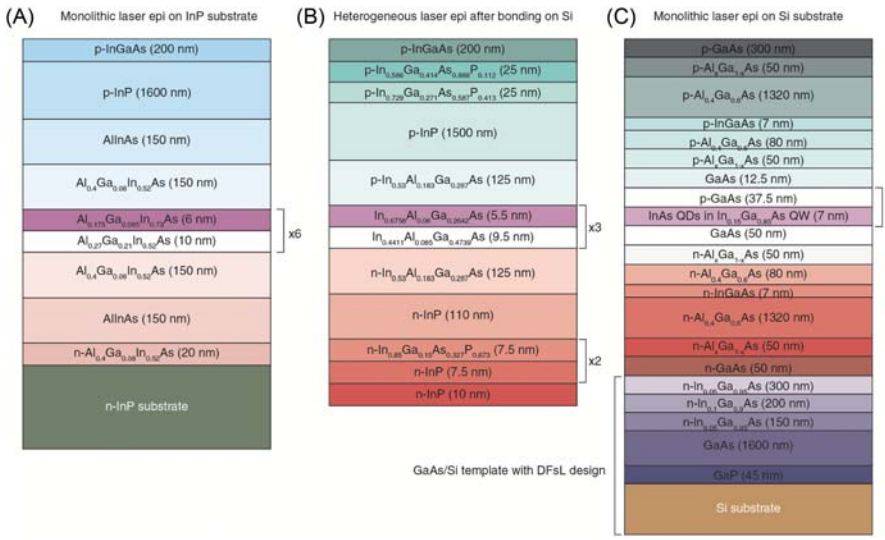

Figure 2.5 Typical epi layer structure comparison for monolithic InP lasers [38], heterogeneous InP lasers on Si [39], and monolithic GaAs lasers on Si [40].

be covered in later sections. Fig. 2.5B shows much better lasing performance of QD-on-Si lasers than that of QW-on-Si lasers in Fig. 2.5A.

2.2 Integration issues

It is widely known that silicon, as an indirect bandgap material, is a poor emitter. Although there are several generations of research toward using silicon or silicon alloy to directly make lasers, their performances are still not sufficient for practical use. Compound III−Vs provide the most efficient lasing and to make integrated lasers with Si or on Si, the key is to integrate III−Vs and Si, from different material groups with disparate properties. In this section, we discuss the issues and solutions of III−V and Si integration.

2.2.1 Process integration

The overall process flow to integrate III−Vs with Si depends on the different integration approaches as discussed in the previous section. A comparison of the overall process timeline and throughput is summarized in Fig. 2.6. For hybrid integration, the III−V gain and Si PICs are processed separately and assembled together. The integration process is die-level and requires accurate optical alignment.

Heterogeneous integration of III−V materials with Si utilizes bonding of III−V films to patterned Si waveguides. The III−V epi wafer can be cleaved into die-size to be bonded onto the active region or be bonded to the Si wafer in full wafer form. The wafer bonding is either based on direct oxygen-assisted plasma-activated bonding or intermediate adhesive BCB bonding. Low surface roughness in both III−V epi and Si is normally required for a high bonding yield in direct bonding. Surface cleanness is also critical to the final bonding yield. Particles left on either the Si wafer or III−V epi could result in bonding failure areas that are over ten times larger than the particle size itself. The laser processing after bonding is thus CMOS-compatible and CMOS facilities can be used for the lithography, etching, doping, passivation, metallization, and so on. Epitaxial growth of III−V QDs on Si substrates has higher requirements than growth on native substrates due to the aforementioned material constraints. To suppress antiphase domains, the silicon wafer used to be (100) with a 6 degree miscut toward (111) [31]. Commercially available germanium on Si wafers was used as growth template in early generations to closely match the lattice constant between Ge and GaAs [32]. Later generations include growth on (100) Si wafers with 4-degree offcut to the (011) plane [33], GaAs-on-V-grooved (001) Si (GoVS) [34], MOCVD (metal organic chemical vapor deposition) grown GaAs/Si substrate [35], and MBE (molecular-beam epitaxy) grown GaP/Si template [36,37], both on (001) Si. On the growth templates with low threading dislocation densities, high-performance lasers are made using a standard laser fabrication process. A comparison of typical laser epi structures is shown in Fig. 2.7.

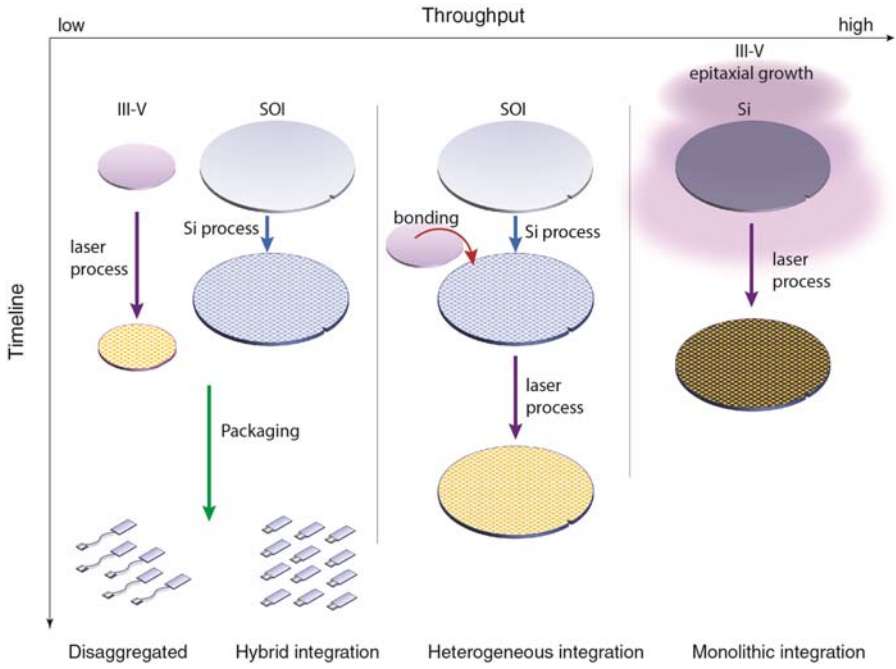

Figure 2.6 A comparison of the process flow for several laser integration methods with Si PICs. Disaggregated lasers and hybrid integration require separate laser and Si processing and the packaging happens at die level. Heterogeneous integration has a wafer bonding process after Si processing and the remaining processes are at wafer-scale. Monolithic integration starts with epitaxial III−V growth on Si and completes with wafer-scale laser processing.

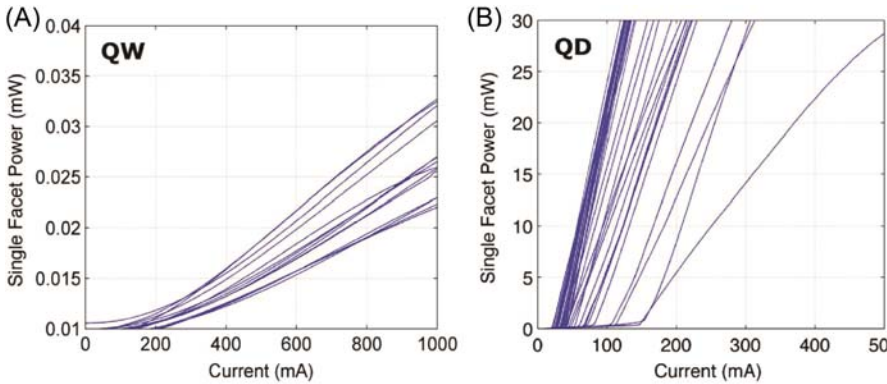

Figure 2.7 A comparison of light-current properties of QW and QD grown on Si lasers. The QW or QD lasers are grown from the same growth template and current−voltage properties [29].

2.2.2 Process temperatures

Process temperature is a key consideration in device fabrication, especially when additional or nonstandard Si IC processing steps are required. Heterogeneous integrated lasers on silicon, due to the mismatched thermal expansion between the III−Vs and Si, due to the bonding process, have thermal budgets that must be carefully defined and managed because excess temperatures will result in bonding or device failure. The bonding temperature varies across different III−V materials. The general trend is that InP which has closer thermal expansion coefficient to Si can allow for a higher bonding temperature than GaAs. In addition, the bonding temperature of a thinner material stack can be higher than that of a thicker stack. For lasers where the passive waveguides (e.g., Si_3N_4) need high-temperature annealing to lower the waveguide loss, the high-temperature process needs to be completed at the front end, before the bonding step.

The MOCVD or MBE growth of III−Vs requires high processing temperatures and multiple thermal cycles to reduce the TDD (threading dislocation density). The fabrication process temperature after the epi growth is lower than the growth temperature so the epitaxial laser thermal budget is generally higher than heterogeneous laser fabrication process.

2.2.3 Yield

The final device yield of integrated lasers for Si PICs is determined by several factors. Hybrid integrated lasers require high-accuracy optical alignment. Robust resistance to vibration is also critical to maintaining a stable laser performance including threshold current, output power, linewidth, and so on. The yield of heterogeneously integrated lasers largely depends on the bonding yield of III−Vs to Si. Close to 100% yield is common in optimized bonding process, which requires correct cleaning procedures, surface activation process, high-temperature annealing, applied pressure, and so on. Due to the bonding of III−V materials and Si that have different thermal expansion coefficients, one concern was that the thermal stress could cause dislocations to form at the bonded interface. However, it was found that the use of strained superlattice layers serves as mechanical barriers to prevent dislocations at the bonding interface from climbing to the active quantum wells. The heterogeneous lasers on silicon show close or even better photoluminescence (PL) performances than lasers grown on native substrates due to dislocations and/or point defects formed during growth gliding away from the active region to the strained superlattice. Other than the bonding process, as a result of the standard semiconductor processing in completing the laser process, the device yield is high and has no fundamental differences from monolithic III−V lasers. The benefit from CMOS facilities is another advantage of laser production for heterogeneous lasers on silicon.

The device yield of heteroepitaxially-grown monolithic lasers on silicon, on the other hand, depends mostly on the growth quality. Any defects present in the laser region could cause the laser to fail, or deteriorate the laser performance and reduce the laser yield. QD active regions show a much higher tolerance to crystalline

defects, compared to QWs. In both cases, cracking in the III−V film can occur and is due to the tensile thermal stress induced during cooling down from the growth temperature. The crack process is well understood [41] and controllable. A laser PL analysis has been done by Liu et al. [31], that compares the InGaAs/GaAs quantum well and InAs/GaAs QD laser material grown on silicon. All the lasers have identical dislocation densities of $2 \times 10^8/cm^2$, identical laser stacks (except QDs or QWs), identical processing and electrical I-V characteristics. The analysis shows that QD grown on Si lasers has only 20% lower PL than QD grown on native GaAs substrates while the PL of QW grown on Si lasers shows over 90% decrease from the native substrate counterparts. For a later study with lower dislocation densities of $7 \times 10^6/cm^2$, the PL of QDs grown on Si lasers is 95% of that of GaAs native counterparts. In addition, while the QD-on-Si lasers show a high device yield at a relative high dislocation density, the QW-on-Si lasers barely lase (see Fig. 2.5A and B comparison).

2.2.4 Reliability requirements

Reliability is a significant consideration for lasers, especially for lasers that are used in uncooled circumstances because many data center applications expose lasers to temperatures as high as 80°C for extended periods of time during their lifespan. For hybrid integrated lasers using separate III−V chip and Si chip, the laser reliability is basically the same as the III−V lasers grown on native substrates, though the mechanical robustness of the coupling between the III−V gain chip and Si coupling interface (grating coupler, edge coupler, photonic wire bonding and so on) is separate and should be considered. On the other hand, the laser reliability is more of an important concern for integrated lasers on silicon since two major different material groups are incorporated in this configuration. With the improvement of the laser epi structure and growth template to minimize dislocation, both heterogeneous lasers on Si and heteroepitaxially grown lasers on Si have seen dramatic improvement in reliability in recent years.

It needs to be noted, heterogeneous-integrated lasers can be integrated with external modulators on the same Si substrate with negligible coupling loss in between, which lowers the requirement for the laser output power. As a result, the drive current can be kept low during operation, which could reduce the junction temperature by several times. Also, the external modulation scheme, unlike direct modulation scheme, does not require a high drive current to increase the extinction ratio and decrease the fall times. In addition, one limiting factor for the reliability of cleaved-facet III−V lasers is that the laser facet could get damaged due to nonradiative recombination of carriers at surface-related defects and the resulting optical absorption can lead to run-away heating [1]. Heterogeneous integrated lasers, on the contrary, could use on-chip Si mirrors to avoid the facet damage and thus reduce both random failure rates and long-term wear out failures.

The superior reliability of heterogeneous InP/Si laser is shown in Fig. 2.8A [23], with 2 failures per billion device hours (2 FIT). In Intel's product, 32 InP/Si-based heterogeneously integrated lasers are used in a 1.6 Tbps photonic engine chiplet design with full laser redundancy to improve PIC reliability [42]. This configuration

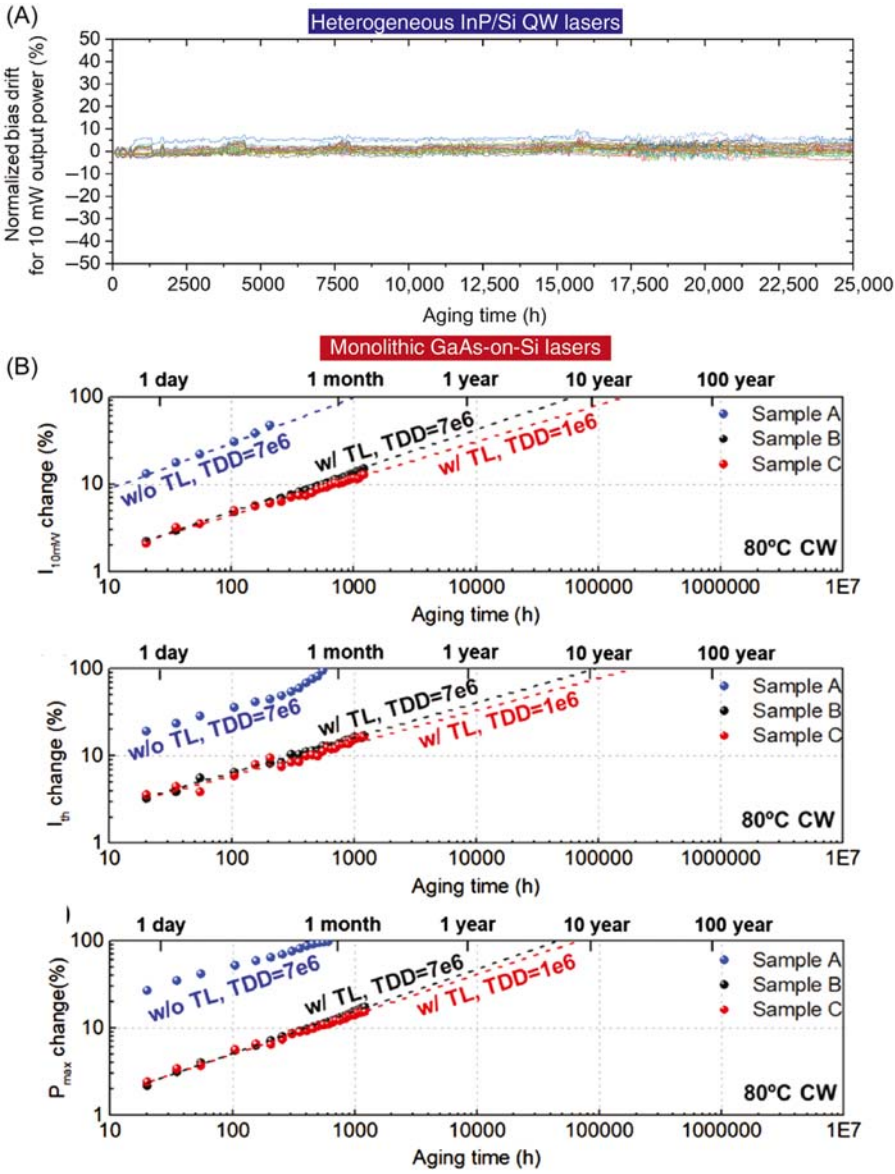

Figure 2.8 Lifetime of heterogeneous InP/Si lasers and monolithic GaAs-on-Si lasers. (A) For 36 III/V-Si lasers operating at a stress environment of 80°C/150 mA over 25,000 hours, the relative drift of bias current for 10-mW laser output power [23]. (B) High-temperature-stable operation of QD-on-Si FP lasers with misfit and dislocation filters. The extracted lifetime at 60°C is over 22 years [40].

is enabled by the dense integration offered by heterogeneous integration. It is difficult to implement with discrete or hybrid-integrated lasers. Aging test of 36 lasers at 80°C with a stress current of 150 mA and for 25,000 hours (Fig. 2.8A) shows the average drift of bias current to produce 10-mW output power is as low as 1%.

For lasers grown on Si, there is a clear advantage of QD active regions over QW, with an improved lifetime of 10^8 hours for QD lasers versus 10^3 for QW. The reduction of threading and misfit dislocation densities plays an important role in improving the performance of QD lasers directly grown on Si. As a result of recent innovations [30], QD laser lifetimes are over 1 million hours at 60°C and 200,000 hours at 80°C. In practical use at data centers, uncooled laser operation is often required, which imposes a strict requirement on the laser high-temperature operation performance. Please see Chapter 13 for additional discussions on quality and reliability.

2.2.5 Cost

The overall cost of lasers has many deciding factors. For instance, packaging normally makes up most of a laser product cost. Traditionally, a complete transceiver requires optical alignment and assembly of discrete optical components including lasers, isolators, modulators, photodetectors, etc., in a transmit optical subassembly and receive optical subassembly. The package assembly has multiple alignment processes and the packaging cost is high. For the laser cost, here, we limit our discussions to the different laser approaches and focus on the cost of materials.

A main cost of semiconductor lasers is the substrate. InP wafers are 2-, 3-, or 4-inch and GaAs wafers recently transited from 3- and 4-inch to 150 mm. Si IC industry uses 150- to 300-mm diameter wafers and there is a clear difference in the scalability of lasers on III−V substrates and lasers on Si substrates. Heteroepitaxial growth on Si provides the lowest-cost solution as it is the only approach that eliminates the expensive InP or GaAs substrates. Meanwhile, heterogeneously integrated lasers or heteroepitaxially grown lasers on Si enable wafer-scale test and burn-in, which also removes a large amount of cost in the full life cycle of a laser product.

Until now, heterogeneous integration is the only approach that allows the wafer-scale integration of lasers with Si PICs including modulators and PDs. The elimination of complex alignments reduced the overall transceiver cost in the packaging stage. As the Si PICs complexity increases, lasers integrated on silicon offer much slower increase rate in the total laser cost compared with discrete laser solutions. By removing the lenses and fiber pigtails that are required for each discrete lasers, heterogeneous-integrated lasers gain additional benefits in the overall cost of a transceiver or photonic engine. To demonstrate the same benefit, heteroepitaxially-grown lasers need to efficiently couple light into passive Si waveguides and other functional photonic devices.

2.2.6 Coupling to Si waveguide

For discrete lasers, the laser output coupling to Si PICs is important in addressing the overall link cost, lane density, power budget, and so on. In many cases, fiber

connections are still required. High efficiency, wide bandwidth, polarization sensitivity, and tolerance in the alignment error are the most common design concerns for different fiber-chip coupling approaches. As the number of signal channels grows, the number of lasers and light coupling pathways also grows, which imposes more challenges in the efficient and uniform coupling and would also ultimately reach limitations from the available chip side density.

The laser out-coupling could use a focusing lens to collimate the laser beam into the optical fiber core or an edge facet coupler of Si PICs. Normally a free-space isolator is required between the lens and the Si chip in a fiber-less configuration. Using fiber connections, a fiber-based isolator is mandatory to minimize the back reflection from the Si chip facet. The need for isolators and back reflection control will be discussed in more detail in the next section. Variations of the direct free space coupling could be found by using 45-degree prism together with collimating prisms to relax the requirements on the fine optical alignment. The laser chip could be flip chip bonded onto a Si PIC and grating couplers are used to facilitate the light coupling into planar Si waveguides.

In fiber-pigtail solutions, the fiber-to-Si PIC coupling mechanism is mainly categorized as edge coupling or vertical coupling. Edge coupling needs to solve mode profile mismatch due to the large dimension mismatch of the silicon waveguide and optical fiber core. Spot-size converters using inverse tapers or deposited low-index SiONx or polymer layers help to enlarge the mode diameter for improved mode overlap with optical fiber mode. Meanwhile, V-groove fiber arrays are utilized to enable higher density of the edge coupling with a fixed pitch size and trench depth for accurate fiber mounting. A different approach is using grating couplers as the out-of-chip diffraction coupler for a large diffraction beam size to match the optical fiber mode. Grating couplers also enable wafer-scale testing without the need to dice or cleave the wafer into dies to access the facet. However, it normally suffers from a relatively narrow coupling efficiency bandwidth ($\sim$40 nm 3-dB bandwidth compared to $>$200 nm for edge coupler). In addition, it has polarization selectivity so in polarization multiplexing (PDM) systems, 2-D grating couplers are normally required to work for both TE and TM modes. Grating coupler-based output couplers can also work with V-groove fiber arrays for a large number of channels.

As mentioned, the introduction of fiber pigtails for disaggregated lasers adds complexity and cost, which will grow significantly with increased number of channels. Especially in PDM systems, polarization-maintaining fibers are required at a much higher cost than single-mode fibers. These issues are completely avoided in heterogeneously integrated lasers as the coupling from lasers to Si waveguides in based on evanescent coupling with the accuracy of CMOS photolithography. This greatly improves the power budget of the entire optical link by removing the coupling loss, and other benefits include less laser driving current which normally helps to increase the WPE, lower the laser junction temperature and increase the laser reliability.

Heteroepitaxially grown lasers on Si substrates need developments in the effective integration with low loss Si waveguides. Progress is being made on heteroepitaxial laser integration on silicon platform and a few other variations. With the advantages of

economy of scale, heteroepitaxially grown epi stack on Si could potentially replace epi stack grown on native substrates for discrete configurations, hybrid integration, or heterogeneous integration through wafer bonding or transfer printing.

2.2.7 Optical stability and feedback tolerance

Lasers are generally sensitive to back-reflection which could destabilize the laser and alter the laser output power and spectral characteristics. Depending on the feedback level, distance to feedback and laser output power, the lasers can operate in multiple regimes that have detrimental consequences, such as broadened linewidth and mode-hopping. As a result, discrete lasers' output is normally followed by an optical isolator in both free space and fiber-based configurations. Lensed fibers with AR coating can be used to reduce the reflection at the fiber-chip coupling interface but isolators are still necessary in most cases for a stable laser operation. For heterogeneous integration, lasers and modulators are on the same chip without intermediate fiber connections or chip facet reflections. Careful designs could minimize the reflections from on-chip Si waveguide components and output facet back to the laser. Error-free operation of InP/Si-based transceiver has been demonstrated without isolators [23].

Another potential solution for the optical feedback-induced instability is to use a laser gain material with a lower linewidth enhancement factor (α-factor). Quantum dot materials exhibit near-zero α-factor with proper p-doping and dot density [43]. Recent work shows promising initial results where penalty-free transmission at 10 GHz under external modulation is achieved with -7.4 dB optical feedback [44]. Low α-factor is also reported with QW-based DMLs utilizing the detuned-loading effect [4]. The α-factor of QW gain materials is reduced to a minimum αeff factor value below 1, for DFB + R laser, two-κ DBR laser, and split-contact DR laser at the optimum mode position. As a result, isolator-free operation with up to 40% reflection is achieved.

2.3 Performance requirements

Data communications impose several requirements on the laser performance, reliability, and cost. A main motivation of moving to silicon photonics for data communications is to achieve a more energy-efficient, reliable, and economically competitive solution. CMOS-compatible processing has unique advantages of high-density integration, device performance, and yield control. Lasers as the essential part of a data link, while having variations based on the architectures and available process capabilities, need to meet several critical performance metrics.

While the need for higher communication bandwidth drives the data rates from 100 to 400 GbE (Gigabit Ethernet) and higher, it also poses more demanding requirements on the lasers. Multiple lasers carrying data at single or multiple wavelengths have been deployed. Next-generation co-packaged optics with higher data

rates (> 12.8 Tb/s) requires a large quantity of laser sources since the lane speed will still be 200 Gbps or less.

2.3.1 Energy requirements

One important motivation for using silicon photonics in data communications is to potentially lower the energy consumption. Lasers normally consume a large portion of the total energy, so they need to improve in energy efficiency. Since the laser standards vary in different transceiver architectures, there is not a standardized metric table but several trends can be drawn.

The power consumption of the laser includes the electrical gain current power, necessary wavelength tuning power, and temperature controlling power if applicable. In data center applications, the lasers are normally working in uncooled operation to eliminate the temperature controlling power. Single wavelength lasers normally have no wavelength tuning power, since the channel guard bands are normally wide enough to tolerate the wavelength drift caused by temperature drifting. So, the majority of laser power consumption results from the laser injection gain current.

The laser WPE characterizes the overall laser energy conversion efficiency. When considering the laser power used in data transmission, the effective electrical-to-optical power conversion factor also depends on the lasing spectrum. For example, multiwavelength lasers may have more comb lines than required in a practical system, which degrades the effective WPE in data transmissions. Taking this into account, we focus on the case of single-wavelength lasers assuming all laser output power (except amplified spontaneous emission) resides in a center wavelength that can be used for data transmission.

Several industry photonics suppliers set the CW-WDM MSA configuration guideline, outlining the laser specifications in grids in the O-band for various applications including co-packaged optics [45]. Based on the CW-WDM MSA Technical Specifications Rev 1.0, each single wavelength laser power is required to be 14 dBm. Fig. 2.9 plots the single wavelength laser power consumption when externally modulated at different speed. It is clearly pointing out that, 1 pJ/bit energy consumption is difficult to achieve unless the single wavelength data rate is exceeding 40 Gb/s/λ. Higher order of modulation, higher modulation speed, and laser fan out into multiple fibers are the solutions to driving down the energy/bit consumption from a single-wavelength laser source.

It is thus critical to use higher data rates to improve the energy efficiency and at the same time several aspects are important. State-of-the art Si modulators provide at most 100 GBaud modulation bandwidth and 100 Gb/s/λ seems to be an upper limit for single wavelength lane with ON-OFF Keying (and without multiplexing or parallelism). Coherent optics still remains at high cost and requires additional power consumption from DSP. More parallelization/multiplexing including pol-mux (polarization multiplexing), PAM-4 modulation, and multiwavelength output from a single laser could be the solutions to further drive down the laser energy consumption.

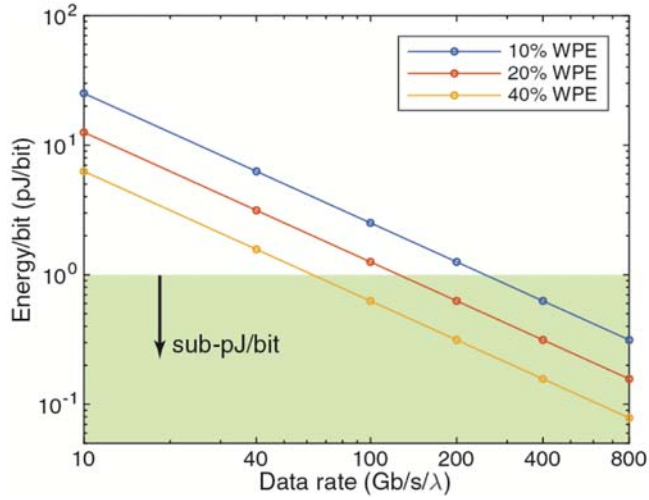

Figure 2.9 Laser source power consumption calculation for different data rates and laser WPE. The calculation is assuming 14 dBm power for single wavelength.

2.3.2 Threshold and slope efficiency

The laser threshold and slope efficiency determine the laser output power and WPE. Also, it sets the laser operation conditions, such as operation current, thus deciding the laser junction temperature, which impacts the laser lifetime. Lasers on Si have demonstrated thresholds below 1 mA [46], but at the expense of output power. QD lasers grown on Si have also demonstrated powers of 175 mW [37] and wall-plug efficiencies of 40%. Generally, a short laser gain section would result in lower threshold current but less peak power at the same time.

2.3.3 Output power

The laser output power requirement varies depending on the transceiver architecture and also the reach of the link. For example, high-power lasers are required in PSM configurations since one laser usually needs to feed multiple channels and accounts for the loss from coupling, splitting, and so on. The typical power requirement for a laser used in PSM4 configuration is over 40 mW for 500 m reach, and over 70 mW for 1 km reach or longer in large-scale data centers.

The laser output power needs to meet the entire link budget. The fiber coupling loss is a major source of additional power requirements, which means the heterogeneous lasers would require less output power if integrated modulators on the Si waveguide circuits are present. In addition, the PD sensitivity which determines the minimum power required for sufficiently low bit error rate is also an important factor in determining the link power budget. High-responsivity photodiodes (PDs) and avalanche PDs could lower the laser power requirements, thus allowing more parallelism lanes and reducing the link energy per bit. However, most transceivers

in data centers need to be MSA (multisource agreement)-compliant so the transmitter output power is predefined.

2.3.4 Spectral characteristics

Compared with FP lasers, high-coherence single longitudinal mode laser is preferred in interconnects with long reach distance and high data rate. For single longitudinal lasers, a high SMSR is required. Multiwavelength lasers (or comb lasers) also require a high extinction ratio in communications systems. Based on the CWDM MSA configuration guideline, the minimum SMSR requirement is 30 dB.

In WDM systems, the laser wavelengths in different channels are set within standardized grids. CWDM in the O-band is currently used as the mainstream technology to allow for process and temperature-induced wavelength drift and to eliminate the individual laser control requirements. The grid spacing in a CWDM4 configuration is 20 nm. With the growing need for more channels and utilization of comb lasers, DWDM is being evaluated for data center interconnects. Depending on the number of channels and wavelength span, the channel spacing could range from 100 to 800 GHz. Fig. 2.10 plots the WDM grids defined by CW-WDM MSA,

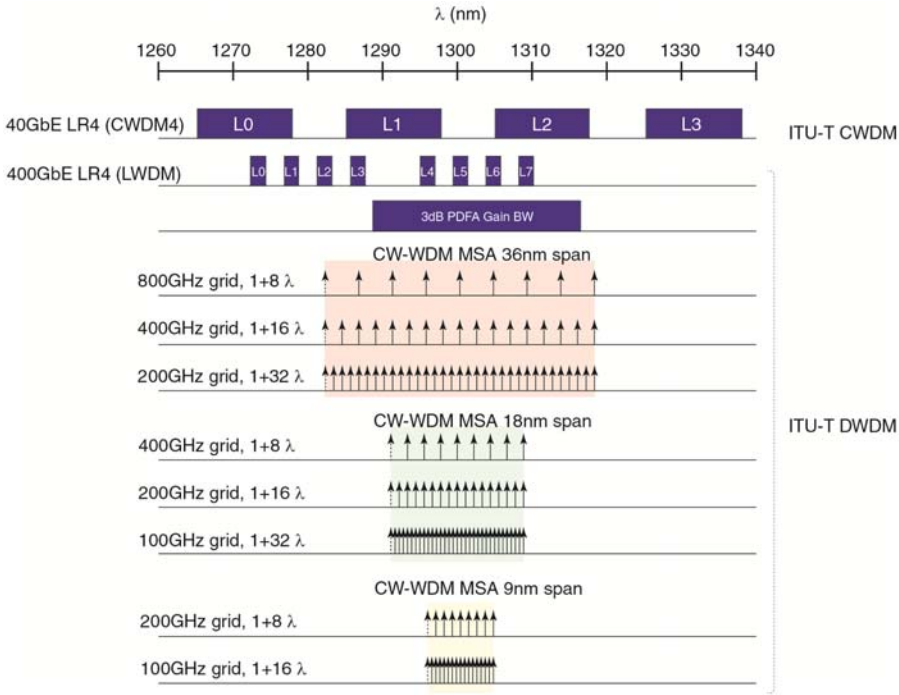

Figure 2.10 CW-WDM MSA defined laser wavelength grids for WDM systems. 8 wavelength grid sets with different span widths (36, 18, and 9 nm) and grid spacing fit into ITU-T DWDM standards. The top columns show the CWDM4 and LWDM grids and 3 dB PDFA Gain bandwidth for a reference [47].

featuring eight wavelength grid sets with three wavelength spans (36, 18, and 9 nm). In each configuration, a clock channel is used for synchronization.

The laser Lorentzian linewidth is required to be better than 20 MHz in the CW-WDM MSA configuration. This requirement is based on the current application situation and minimum signal processing requirements, that is, no high-order modulation formats. A narrow laser linewidth benefits a long transmission distance and high-order modulation formats with potential DSP-free operation, which is important for future data interconnects when the coherent communication is introduced in data center interconnects. An estimation of the laser linewidth requirement for different formats is calculated in [48]. Sub-kHz laser linewidth is required for more advanced formats than 64-QAM. In a later section, the recent progress of narrow-linewidth lasers using hybrid integration and heterogeneous integration will be introduced.

2.3.5 Intensity noise

Intensity noise from the laser could result in errors at the receiver side. To lower the intensity noise, lasers' power could be biased at higher current to lower the RIN, thus introducing additional laser power penalty. A laser RIN below −135 dBc/Hz is required in the CW-WDM MSA configuration.

2.3.6 Operation temperature

The data center normally operates at a high temperature and reliable operation up to 80°C is required. The lasers are normally under uncooled operation to save the energy use, which imposes a strict requirement on the laser high-temperature performance. With high operation temperature, the laser efficiency drops and reliability becomes more questionable. Normally the lasers are required to perform per specification at a temperature ranging from 20°C to 80°C. In WDM systems, the temperature-dependent wavelength shift of DFB lasers is on the order of 0.1 nm/°C while the III−V gain peak moves at around 0.5 nm/°C. For a 60°C temperature change, the shift in gain peak is thus around 30 nm, which makes the relation between DFB lasing wavelength and gain peak wavelength offset an important design parameter in uncooled DFB lasers for high-efficient lasing across the full temperature range.

InP-based lasers are currently dominating the laser market for datacom and telecom. The mature quaternary alloys InGaAsP or InGaAlAs are lattice matched to InP and both can produce efficient QW-based lasers. The InGaAlAs material system has advantages for high temperatures operation due to its larger conduction band offset, which in general provides better carrier confinement and enhances the high-temperature operation performance.

With an atom-like discrete density of states, QDs show promising high-temperature performance thanks to the large quantized-energy separation that prevents carrier population broadening into higher states at elevated temperatures. As a result, QD lasers are reported to lase at a temperature as high as 220°C [49], exceeding the

best-performing QW counterparts (around 185°C for InGaAlAs QW lasers [50]). In addition, the characteristic temperature T_0 in determining the lasing threshold $I_{th} = I_0 e^{T/T0}$ for QD is higher than that of QW, as QD lasers are less sensitive to temperature variation in practical data center operations.

2.4 State-of-the-art

This section will review the recent progress of state-of-the-art silicon photonic compatible lasers for data center applications. As discussed before, we will cover both disaggregated lasers and integrated lasers with Si.

2.4.1 Disaggregated

Due to the present difficulty of integrating III−V with Si, disaggregated lasers are being explored as a solution for both pluggable optics modules and co-packaged optics. In contrast to the pluggable optical transceivers where disaggregated lasers can be packaged with Si PICs through free-space lens-based coupling, in co-packaged optics, a laser supply with multiple laser arrays is remotely placed away from the central switch core. Thus, fiber connections are required between the laser and Si PICs.

InP-based lasers are a mature technology and the progress is mostly focused on the laser architectures. For co-packaged optics, the laser supply could remain at the switch box, in the form of pluggable modules which allow replacement upon malfunction or failure. Fiber connections are used to connect the laser output to the Si PICs. Laser supplies are required to power more than 1 laser in CPO architectures, which requires high output power of individual lasers, depending on the overall link budget and source of loss.

2.4.2 Integrated

Integrated lasers with Si PICs offer the advantages of utilizing the low loss Si waveguide to achieve laser performance that is unattainable with monolithic III−Vs. A prime example is the laser linewidth reduction progression for different laser integration approaches as shown in Fig. 2.11. Meanwhile, as the laser is directly integrated with Si waveguide, it removes the requirement for additional fiber coupling to improve the power budget. There has been tremendous progress in aforementioned integration approaches for laser gain with Si PICs.

2.4.2.1 Hybrid

Hybrid integration provides a laser integration solution to combine the benefits of III−V gain materials and Si-based passive materials at the same time. Since integration is based on the light coupling between separate chips, both III−V gain chip and passive Si-chip can be optimized separately without the limitation of

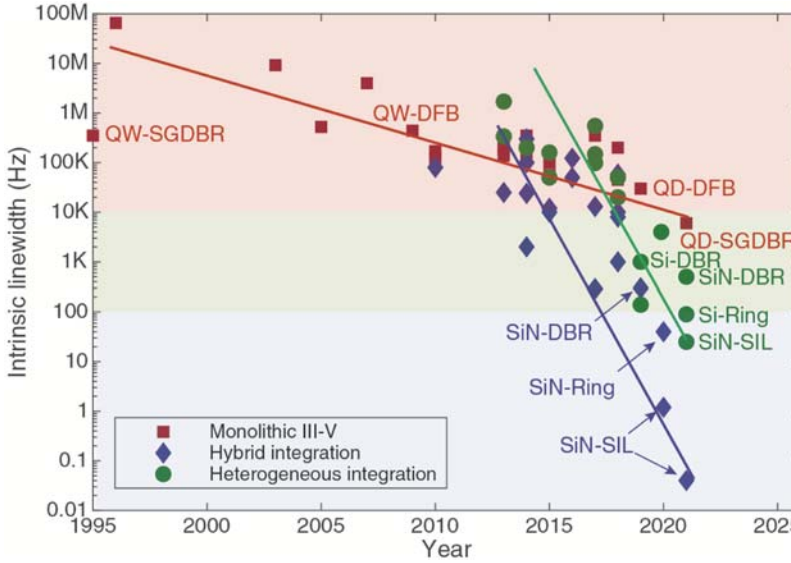

Figure 2.11 Semiconductor laser linewidth progression. The fundamental linewidth progress for different integration systems and operation principles [51].

integration compatibility issues. III−V gain chip at different wavelength bands can be used depending on the applications, and are placed in Si-based passive cavities of different designs. With further reduced optical loss, Si_3N_4-based cavities are finding potential for further improving the performance [52].

Tunable lasers based on ring resonators are examples of Si- or Si_3N_4-based cavities offering advantages in laser linewidth over III−V-based tunable lasers due to the low waveguide losses. Kobayashi et al. demonstrated Si-photonic hybrid ring external cavity-based wavelength tunable lasers. The device is based on passively aligned gain chip and Si ring-resonator-based reflector chip. The laser output from the gain chip is coupled to another booster SOA. With Si waveguide loss lower than 0.5 dB/cm, the obtained laser linewidth is below 15 kHz across the full C band. The laser full tuning range is 65 nm, and the laser output power after SOA is over 100 mW [16]. Later research showcases that a hybrid-integrated tunable laser with 60 kHz laser linewidth, 21.5 dBm fiber coupled output power could achieve accurate laser wavelength tuning and locking. Using this narrow-linewidth laser, coherent optical transmission using 64 Gbaud 64-QAM data transmission is successfully demonstrated [53].

By replacing Si with Si_3N_4, the waveguide loss can be further reduced by up to three orders of magnitude. State-of-the-art Si_3N_4 waveguide loss can be lower than 0.1 dB/m through proper geometry engineering and high-temperature annealing [11]. Using hybrid integration, Si_3N_4-based ring resonators could enable a longer external cavity length without significantly increasing the cavity loss. As a result, Si_3N_4-based hybrid integrated laser achieved fundamental linewidth as low as

40 Hz across a wide tuning range [54]. Single-wavelength lasers using extended Si_3N_4 Bragg grating also result in around 320 Hz Lorentzian linewidth, offering a competitive configuration in DBR solutions for laser arrays [55].

Hybrid laser integration is also starting to be introduced in silicon photonics foundry. For example, GlobalFoundries Inc. introduced an approach to flip chip bonding a III−V laser in an etched Si trench that is formed together with coupling fiber V-groove arrays [19]. The laser-PIC coupling is achieved by precise mechanical stops and optical alignment features. Up to 11 dBm Si waveguide output power is achieved and the laser shows no sign of degradation in RIN and mode hop performance compared to free running III−V lasers. Sub-μm misalignment could ensure that the coupling loss is below around 1 dB.

Photonic wire bonding (PWB) is a new technology in addressing the laser-Si PIC coupling bottleneck. Using direct-write two-photon lithography for three-dimensional freeform waveguides fabrication, coupling losses below 0.4 dB are achieved between InP laser chip and Si PIC [18]. Using the same PWB technology, hybrid InP/Si lasers based on Si external cavities are also demonstrated [56]. PWB, as a replacement for free space coupling, enables low-loss and flexible hybrid integration of chips made from different material systems. The stability of the PWB against environmental vibrations is still a concern and is under further investigation.

As mentioned before, hybrid integration provides a method to adopt the best-in-class actives and passives in an assembly for optimized laser performance. Utilizing efficient gain from high-power InP lasers and high-Q from ultralow loss Si_3N_4 waveguides, SIL could enable ultralow-noise lasers and OFCs [11,14]. This scheme shows that integrated laser performances can be comparable with fiber lasers and the ultranarrow laser linewidth is promising for next-generation optical data communication links.

2.4.2.2 Heterogeneous

Heterogeneous integration evolves rapidly in both research and commercially available product developments, progressing from proof-of-concept demonstrations to industry products shipping millions per year for data center interconnects [23]. Recently, Tower Semiconductor Ltd. and Intel Foundry Services also announced the availability of heterogeneous lasers in coming multiproject wafer schedules. Heterogeneous-integrated lasers nowadays not only contain a complete list of laser types including single wavelength lasers (DFBs, DBRs) and multiwavelength lasers (mode-locked lasers) but also exhibit fundamental superiority in laser linewidth compared to monolithic III−V disaggregated lasers. Several demonstrations highlighting the narrow-linewidth heterogeneous lasers are shown in Fig. 2.12. In addition, heterogeneous integration enables novel devices, such as laser soliton microcomb devices, that are not achievable with monolithic III−V integration.

Heterogeneously integrated DFB laser arrays for CWDM systems are used in Intel products with superior performance including output power and reliability. The lasers can last up to 150°C and at 80°C less than 100 mA still produces 20 mW output power [23]. As mentioned before, the low-current-density operation and

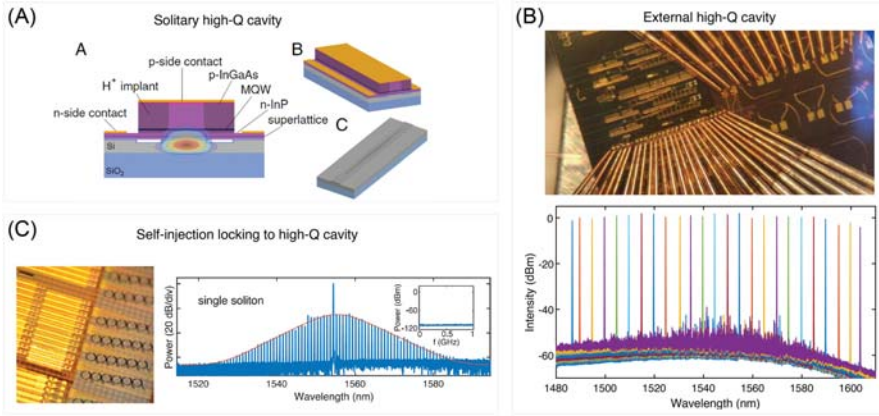

Figure 2.12 A comparison of approaches to achieve narrow linewidth heterogeneous III−V/Si lasers. (A) Solitary laser using high-Q Si cavity [57]. (B) External-cavity tunable laser using cascaded high-Q Si ring resonators [57]. (C) Self-injection locking of low-coherence DFB lasers to high-Q Si_3N_4 ring resonators [12].

associated low junction temperature improve laser stability. The wafer-scale integration of lasers is well suited for scaling up in configurations where more laser channels are required to reduce the energy per bit. In addition, the Intel approach employs one backup laser per channel followed by an MZI-based switch and results in further improved reliability over 2 failures per billion device hours (2 FIT) achieved per single laser [23].

For WDM systems, multiwavelength sources are attractive solutions to convey wavelength multiplexed data, including mode-locked comb lasers and optical Kerr combs. Combined with MRM arrays, they remove the need for spectral multiplexers and demultiplexers and enable small device footprint. Replacing DFB arrays with a single comb-wavelength laser also removes the need to control the channel spacing between individual lasers. In addition, since the intrinsic energy consumed by a laser to reach threshold does not scale with the number of channels for comb lasers, multiwavelength lasers could lower the total power consumption for WDM sources, compared with using single-wavelength laser arrays. Heterogeneous integrated mode-locked lasers are seeing tremendous progress in addressing the power consumption and improving the noise performance, with low loss Si waveguide integration [58].

Tunable lasers using heterogeneous III−V/Si integration could have similar benefits as hybrid integration including low-loss external cavities. Thanks to the evanescent coupling in heterogeneous integration, the laser gain and external cavity have little transition loss and reflection, could eliminate the hybrid packaging constraints and enable more laser channels directly on chip. Si ring resonators integration with InP QW or GaAs QD gain section enables C/O-band lasers with record sub-100 Hz fundamental linewidth or several kHz depending on the resonator cavity designs. With carefully designed ring parameters, three- or four-ring resonators

could extend the laser tuning range to be over 110 nm, covering S, C, and L bands simultaneously [8,59].

Heterogeneous integration also achieves significant progress in recent years in integrating low-index Si_3N_4 waveguides. The large index mismatch between III−V epi and Si_3N_4 does not allow direct III−V/Si_3N_4 wafer bonded structure for efficient evanescent coupling between the active and passive layers. So, a multilayer structure using Si as an intermediate layer is developed to improve the mode transition efficiency between the active III−V and passive Si_3N_4 layers [39,60]. The lowest linewidth achieved is below 400 Hz and the output power in the Si_3N_4 waveguide is over 25 mW. These lasers, due to the presence of Si layer, could be further integrated with mature Si photonic devices including modulators and PDs for complex PICs. Noteworthily, due to the low thermo-optic coefficient of Si_3N_4 and SiO_2, the lasing wavelength is stable and its drift rate over temperature is only 10 pm/°C. This low-temperature dependence can be combined with Si_3N_4 AWGs for "athermal" laser architecture designs in future transceivers and potentially enable DWDM applications due to the less thermal-induced channel crosstalk. Another benefit of using Si_3N_4 is to avoid the nonlinear loss that could limit the power handling capability of Si and be detrimental to tunable lasers performance [61].

The ability to integrate efficient III−V gain and ultra-low-loss waveguides gave rise to the fully integration of lasers and nonlinear high-Q ring resonators. Traditionally, soliton microcomb generation requires high-power table-top pump lasers, amplifiers, and fiber connections. Although soliton microcomb could provide a large number of carriers from on-chip resonators, the lasers were discrete and not integrated. Recently demonstrated laser soliton microcombs combine InP/Si DFB lasers and nonlinear high-Q Si_3N_4 ring resonators on a monolithic Si substrate [12]. Using electrically controlled laser gain current and phase current, the device output 100-GHz repetition rate soliton states. The comb lines show reduced laser frequency noise due to the SIL. This heterogeneous device provides a type of integrated device that is potentially applicable for future ultra-high-capacity data interconnects. Before practical applications, work is needed to improve the pump-comb power conversion efficiency and spectral flatness across the comb lines.

2.4.2.3 Monolithic

The last decade not only witnessed the first demonstrations of monolithic growth of QDs materials on Si but also the tremendous progress in this field. The improvements include material growth and epitaxial layer designs. They improved the laser performance in power, threshold, slope efficiency, high-temperature operation, and long-term reliability. With optimized buffer designs, low TDD down to around $1 \times 10^6/cm^2$ has been demonstrated using strained layer superlattices and thermal anneals [41]. In addition, the misfit dislocations in active layers due to the propagation of threading dislocations during cooldown could be minimized by adding blocking layers on either side of the active layer. High-temperature QD lasers on Si thus have an increased lifetime of around 22 years at elevated 80°C ambient temperature, fulfilling the requirements of data center applications. More work is being

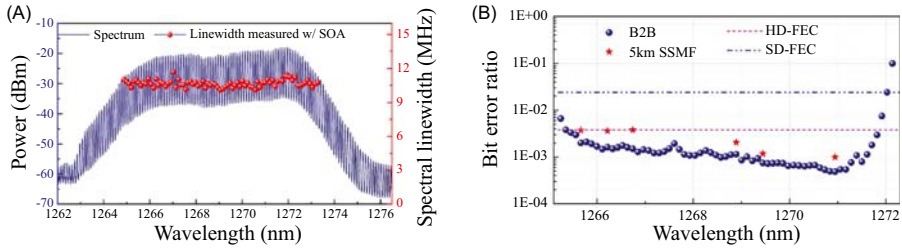

Figure 2.13 Quantum dot lasers on Si for data communications. (A) High-channel count QD-on-Si mode-locked lasers for (B) 4.1 Tb/s using 65 channels for data communication [62].

done to investigate the laser growth in SOI trenches to effectively couple laser output into SOI waveguides.

QD-based mode-locked lasers benefit from the inhomogeneous broadening gain. Fig. 2.13 shows 4.1 Tb/s transmission using 64 wavelength channels from a 20 GHz QD-on-Si mode-locked-laser [62] showing the promise of compact sources in very high-capacity systems. Recent developments on QD material growth, improving lasers' performance, and laser integration with Si PICs are summarized in detail [30,63].

2.5 Outlook

2.5.1 Next 5 years

Silicon PICs and data center interconnects are advancing at a high pace in terms of the integration level (Fig. 2.14) and feeding the growing switch bandwidth (Fig. 2.15). For the next 5 years, silicon PICs will be widely deployed with both disaggregated remote laser sources and integrated laser sources from different suppliers. Multiwavelength lasers for WDM architecture with more channels and higher speed lane will be deployed. Coherent systems in datacom will be emerged as one option in data centers when the performance/cost advantage is validated. Lasers for coherent systems have more requirements on the performance including tunability, linewidth, RIN, etc., and modern lasers on silicon are a promising option to meet these requirements

High-power disaggregated laser sources to feed more channels are already an urgent issue, especially for CPO architecture. Due to the coupling loss and channel MUX/deMUX loss, the requirement for disaggregated laser output power is expected to be $>$20 dBm per wavelength. The laser reliability is becoming a concern at high current. At the same time, DMLs are still evolving with new methods to enable up to 200 GbE per lane.

Heterogeneous-integrated lasers will become more prevalent as the number of wavelengths, number of spatial channels, and general complexity increase at higher capacities. The ability to self-test on wafer, 3D bonding, etc., drives the need for

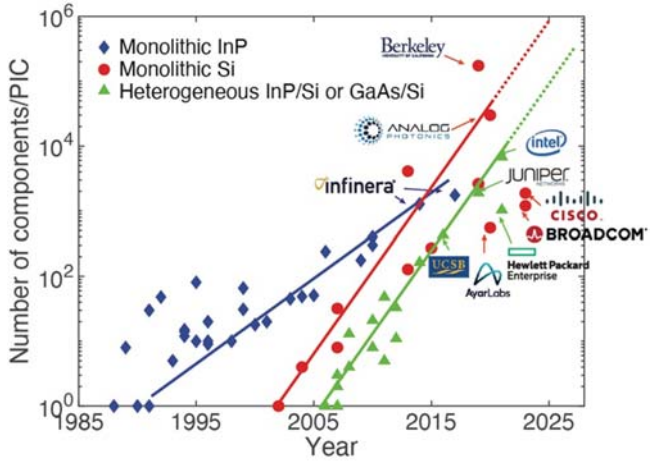

Figure 2.14 For three photonic integration platforms: InP substrate with integrated lasers, Si substrate without integrated lasers, and Si substrate with integrated lasers, the number of photonic components integrated on a single waveguide increased over time [51].

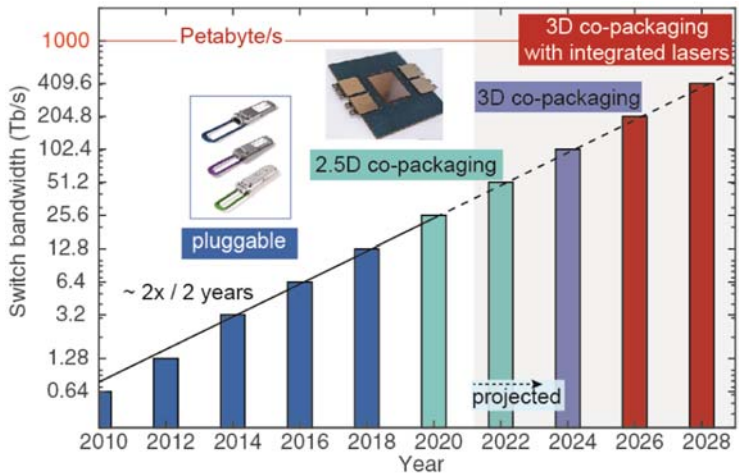

Figure 2.15 The evolution and future of switch bandwidth over four generations. Inset images show pluggable transceivers from Intel and a 2.5D co-packaged PICs and switch from Broadcom Inc. The year denotes when the technology was first used [51].

integration. The unique advantage of direct on-chip integration also provides additional flexibility to deploy redundant lasers to reduce the overall failure rate by orders of magnitude. Challenges include the operation temperature of the entire PIC in CPO configurations and temperature control of lasers, modulators, etc.

2.5.2 Next 20 years

For a longer time span, integrated lasers with Si PICs are necessary to support more parallel channels at low coupling loss and relieve the strict requirements on laser output power. The adoption of more lasers instead of splitting power from single laser also avoids the nonlinear loss limited power handling bottleneck of Si. The integration will be heterogeneous or heteroepitaxial based on the technology availability and the market demands.

DWDM system using comb sources including mode-locked lasers and Kerr combs is required as they can meet the power efficiency requirements and channels requirements. Conversion efficiency of Kerr combs needs to be improved and the pump laser integration will keep evolving.

Monolithic integration of lasers on Si will be an important option to realize the lowest-cost lasers with dense integration with Si PICs. Heteroepitaxial growth of QD lasers on SOI wafers or Si wafers with SiN waveguides is the key to transforming lasers grown on Si to lasers applicable for advanced Si PICs.

References

[1] L.A. Coldren, S.W. Corzine, M.L. Mashanovitch, Diode Lasers and Photonic Integrated Circuits, vol. 218, John Wiley & Sons, 2012.
[2] A. Liu, P. Wolf, J.A. Lott, D. Bimberg, Vertical-cavity surface-emitting lasers for data communication and sensing, Photonics Res. 7 (2) (2019) 121−136.
[3] D. Liang, A. Roshan-Zamir, Y.H. Fan, C. Zhang, B. Wang, A. Descos, et al., Fully-integrated heterogeneous dml transmitters for high-performance computing, J. Light Technol. 38 (13) (2020) 3322−3337.
[4] Y. Matsui, R. Schatz, D. Che, F. Khan, M. Kwakernaak, T. Sudo, Low-chirp isolator-free 65-GHz-bandwidth directly modulated lasers, Nat. Photonics 15 (1) (2021) 59−63.
[5] P.A. Morton, M.J. Morton, High-power, ultra-low noise hybrid lasers for microwave photonics and optical sensing, J. Light. Technol. 36 (21) (2018) 5048−5057.
[6] L.A. Coldren, G. Alan Fish, Y. Akulova, J.S. Barton, L. Johansson, C.W. Coldren, Tunable semiconductor lasers: a tutorial, J. Light Technol. 22 (1) (2004) 193.
[7] B. Liu, A. Shakouri, J.E. Bowers, Wide tunable double ring resonator coupled lasers, IEEE Photonics Technol. Lett. 14 (5) (2002) 600−602.
[8] M.A. Tran, D. Huang, J.E. Bowers, Tutorial on narrow linewidth tunable semiconductor lasers using Si/III-V heterogeneous integration, APL Photonics 4 (11) (2019) 111101.
[9] W.W. Chow, S. Liu, Z. Zhang, J.E. Bowers, M. Sargent, Multimode description of self-mode locking in a single-section quantum-dot laser, Opt. Express 28 (4) (2020) 5317−5330.
[10] B. Stern, X. Ji, Y. Okawachi, A.L. Gaeta, M. Lipson, Battery-operated integrated frequency comb generator, Nature 562 (7727) (2018) 401−405.
[11] W. Jin, Q.F. Yang, L. Chang, B. Shen, H. Wang, M.A. Leal, et al., Hertz-linewidth semiconductor lasers using CMOS-ready ultra-high-Q microresonators, Nat. Photonics 15 (5) (2021) 346−353.
[12] C. Xiang, J. Liu, J. Guo, L. Chang, R.N. Wang, W. Weng, et al., Laser soliton microcombs heterogeneously integrated on silicon, Science 373 (6550) (2021) 99−103.

[13] J.K. Jang, Y. Okawachi, Y. Zhao, X. Ji, C. Joshi, M. Lipson, et al., Conversion efficiency of soliton Kerr combs, Opt. Lett. 46 (15) (2021) 3657−3660.
[14] B. Shen, L. Chang, J. Liu, H. Wang, Q.F. Yang, C. Xiang, et al., Integrated turnkey soliton microcombs, Nature 582 (7812) (2020) 365−369.
[15] Y. Matsui, Directly modulated laser technology: past, present, and future, in: Datacenter Connectivity Technologies: Principles and Practice, River, 2018, pp. 87−211.
[16] N. Kobayashi, K. Sato, M. Namiwaka, K. Yamamoto, S. Watanabe, T. Kita, Hirohito Yamada, and Hiroyuki Yamazaki. Silicon photonic hybrid ring-filter external cavity wavelength tunable lasers, J. Light. Technol. 33 (6) (2015) 1241−1246.
[17] B. Song, C. Stagarescu, S. Ristic, A. Behfar, J. Klamkin, 3D integrated hybrid silicon laser, Opt. Express 24 (10) (2016) 10435−10444.
[18] M.R. Billah, M. Blaicher, T. Hoose, P.I. Dietrich, P. Marin-Palomo, N. Lindenmann, et al., Hybrid integration of silicon photonics circuits and InP lasers by photonic wire bonding, Optica 5 (7) (2018) 876−883.
[19] Y. Bian, K. Ramachandran, B. Peng, B. Hedrick, S. Mills, K. Donegan, et al., Hybrid III-V laser integration on a monolithic silicon photonic platform, in: Optical Fiber Communication Conference, Optical Society of America, 2021, pp. M5A−2.
[20] M. Mack, A. Ayazi, Y. Chi, A. Dahl, P. De Dobbelaere, S. Denton, et al., Luxtera's silicon photonics platform for transceiver manufacturing, in: Proceedings of the International Conference on Solid State Devices and Materials, 2014, pp. 506−507.
[21] B. Li, W. Jin, L. Wu, L. Chang, H. Wang, B. Shen, et al., Reaching fiber-laser coherence in integrated photonics, Opt. Lett. 46 (20) (2021) 5201−5204.
[22] R. Blum, Integrated silicon photonics for high-volume data center applications, in: Optical Interconnects XX, vol. 11286, International Society for Optics and Photonics, 2020, p. 112860M.
[23] R. Jones, P. Doussiere, J.B. Driscoll, W. Lin, H. Yu, Y. Akulova, et al., Heterogeneously integrated InP/silicon photonics: fabricating fully functional transceivers, IEEE Nanotechnol. Mag. 13 (2) (2019) 17−26.
[24] C. Xiang, W. Jin, D. Huang, M.A. Tran, J. Guo, Y. Wan, et al., High-performance silicon photonics using heterogeneous integration, IEEE J. Sel. Top. Quantum Electron. 28 (3) (2022) 1−15.
[25] M.L. Davenport, S. Skendžić, N. Volet, J.C. Hulme, M.J.R. Heck, J.E. Bowers, Heterogeneous silicon/III-V semiconductor optical amplifiers, IEEE J. Sel. Top. Quantum Electron. 22 (6) (2016) 78−88.
[26] Y. Wan, C. Xiang, J. Guo, R. Koscica, M.J. Kennedy, J. Selvidge, et al., High speed evanescent quantum-dot lasers on Si, Laser Photonics Rev. 15 (8) (2021) 2100057.
[27] T. Hiraki, T. Aihara, K. Hasebe, K. Takeda, T. Fujii, T. Kakitsuka, et al., Heterogeneously integrated III-V/Si MOS capacitor Mach-Zehnder modulator, Nat. Photonics 11 (8) (2017) 482−485.
[28] X. Xie, Q. Zhou, E. Norberg, M. Jacob-Mitos, Y. Chen, A. Ramaswamy, et al., Heterogeneously integrated waveguide-coupled photodiodes on soi with 12 dBm output power at 40 GHz, in: Proceedings of the Optical Fiber Communication Conference, Optical Society of America, 2015, pp. Th5B–7.
[29] A.Y. Liu, J. Bowers, Photonic integration with epitaxial III-V on silicon, IEEE J. Sel. Top. Quantum Electron. 24 (6) (2018) 1−12.
[30] C. Shang, Y. Wan, J. Selvidge, E. Hughes, R. Herrick, K. Mukherjee, et al., Perspectives on advances in quantum dot lasers and integration with Si photonic integrated circuits, ACS Photonics 8 (9) (2021) 2555−2566.

[31] A.Y. Liu, S. Srinivasan, J. Norman, A.C. Gossard, J.E. Bowers, Quantum dot lasers for silicon photonics, Photonics Res. 3 (5) (2015) B1−B9.
[32] A.Y. Liu, C. Zhang, J. Norman, A. Snyder, D. Lubyshev, J.M. Fastenau, et al., High performance continuous wave 1.3 μm quantum dot lasers on silicon, Appl. Phys. Lett. 104 (4) (2014) 041104.
[33] S. Chen, W. Li, J.W. Qi, J. Mingchu, T. Samuel Shutts, S.N. Elliott, et al., Electrically pumped continuous-wave III-V quantum dot lasers on silicon, Nat. Photonics 10 (5) (2016) 307−311.
[34] J. Norman, M.J. Kennedy, J. Selvidge, Q. Li, Y.W. Alan, Y. Liu, et al., Electrically pumped continuous wave quantum dot lasers epitaxially grown on patterned, on-axis (001) Si, Opt. Express 25 (4) (2017) 3927−3934.
[35] S. Chen, M. Liao, M. Tang, J. Wu, M. Martin, T. Baron, et al., Electrically pumped continuous-wave 1.3 μm InAs/GaAs quantum dot lasers monolithically grown on on-axis Si (001) substrates, Opt. Express 25 (5) (2017) 4632−4639.
[36] A.Y. Liu, J. Peters, X. Huang, D. Jung, J. Norman, M.L. Lee, et al., Electrically pumped continuous-wave 1.3 μm quantum-dot lasers epitaxially grown on on-axis GaP/Si, Opt. Lett. 42 (2) (2017) 338−341.
[37] D. Jung, J. Norman, M.J. Kennedy, C. Shang, B. Shin, Y. Wan, et al., High efficiency low threshold current 1.3 μm inas quantum dot lasers on on-axis (001) GaP/Si, Appl. Phys. Lett. 111 (12) (2017) 122107.
[38] J. Piprek, J.K. White, A.J. Spring Thorpe, What limits the maximum output power of long-wavelength AlGaInAs/InP laser diodes? IEEE J. Quantum Electron. 38 (9) (2002) 1253−1259.
[39] C. Xiang, W. Jin, J. Guo, J.D. Peters, M.J. Kennedy, J. Selvidge, et al., Narrow-linewidth III-V/Si/Si_3N_4 laser using multilayer heterogeneous integration, Optica 7 (1) (2020) 20−21.
[40] C. Shang, E. Hughes, Y. Wan, M. Dumont, R. Koscica, J. Selvidge, et al., High-temperature reliable quantum-dot lasers on si with misfit and threading dislocation filters, Optica 8 (5) (2021) 749−754.
[41] C. Shang, J. Selvidge, E. Hughes, J.C. Norman, A.A. Taylor, A.C. Gossard, et al., A pathway to thin gaas virtual substrate on on-axis Si (001) with ultralow threading dislocation density, Phys. Status Solidi (a) 218 (3) (2021) 2000402.
[42] S. Fathololoumi, D. Hui, S. Jadhav, J. Chen, K. Nguyen, M.N. Sakib, et al., 1.6 Tbps silicon photonics integrated circuit and 800 Gbps photonic engine for switch co-packaging demonstration, J. Light. Technol. 39 (4) (2021) 1155−1161.
[43] Z. Zhang, D. Jung, J.C. Norman, W.W. Chow, J.E. Bowers, Linewidth enhancement factor in InAs/GaAs quantum dot lasers and its implication in isolator-free and narrow linewidth applications, IEEE J. Sel. Top. Quantum Electron. 25 (6) (2019) 1−9.
[44] J. Duan, H. Huang, B. Dong, D. Jung, J.C. Norman, J.E. Bowers, et al., 1.3-μm reflection insensitive InAs/GaAs quantum dot lasers directly grown on silicon, IEEE Photonics Technol. Lett. 31 (5) (2019) 345−348.
[45] M. Sysak, J. Johnson, D. Lewis, C. Cole, CW-WDM MSA Technical Specifications Rev 1.0. 2021.
[46] Y. Wan, J. Norman, Q. Li, M.J. Kennedy, D. Liang, C. Zhang, et al., 1.3 μm submilliamp threshold quantum dot micro-lasers on si, Optica 4 (8) (2017) 940−944.
[47] C. Cole, A new specification for multi-wavelength optical laser sources for advanced integrated optics, in: Proceedings of the European Conference on Optical Communication (ECOC), 2021.
[48] M., Seimetz, Laser linewidth limitations for optical systems with high-order modulation employing feed forward digital carrier phase estimation, in: Proceedings of the OFC/NFOEC

2008-2008 Conference on Optical Fiber Communication/National Fiber Optic Engineers Conference, IEEE, 2008, pp. 1−3.
[49] T. Kageyama, K. Nishi, M. Yamaguchi, R. Mochida, Y. Maeda, K. Takemasa, et al., Extremely high temperature (220 °C) continuous-wave operation of 1300-nm-range quantum-dot lasers, in: Proceedings of the European Conference on Lasers and Electro-Optics, Optical Society of America, 2011, p. PDA 1.
[50] C.E. Zah, R. Bhat, T.P. Lee, High temperature operation of AlGaInAs/InP lasers, in: Proceedings of the Seventh International Conference on Indium Phosphide and Related Materials, IEEE, 1995, pp. 14−17.
[51] N. Margalit, C. Xiang, S.M. Bowers, A. Bjorlin, R. Blum, J.E. Bowers, Perspective on the future of silicon photonics and electronics, Appl. Phys. Lett. 118 (22) (2021) 220501.
[52] D. Dai, J. Bauters, J.E. Bowers, Passive technologies for future large-scale photonic integrated circuits on silicon: polarization handling, light non-reciprocity and loss reduction, Light: Sci. Appl. 1 (3) (2012) e1.
[53] Y. Gao, J.C. Lo, S. Lee, R. Patel, L. Zhu, J. Nee, et al., High-power, narrow-linewidth, miniaturized silicon photonic tunable laser with accurate frequency control, J. Light. Technol. 38 (2) (2020) 265−271.
[54] Y. Fan, A. Van Rees, P.J.M. Van Der Slot, J. Mak, R.M. Oldenbeuving, M. Hoekman, et al., Hybrid integrated InP-Si_3N_4 diode laser with a 40-Hz intrinsic linewidth, Opt. Express 28 (15) (2020) 21713−21728.
[55] C. Xiang, P.A. Morton, J.E. Bowers, Ultra-narrow linewidth laser based on a semiconductor gain chip and extended Si_3N_4 Bragg grating, Opt. Lett. 44 (15) (2019) 3825−3828.
[56] Y. Xu, P. Maier, M. Blaicher, P.I. Dietrich, P.M. Palomo, W. Hartmann, et al., Hybrid external-cavity lasers (ECL) using photonic wire bonds as coupling elements, Sci. Rep. 11 (1) (2021) 1−10.
[57] C.T. Santis, S.T. Steger, Y. Vilenchik, A. Vasilyev, A. Yariv, High-coherence semiconductor lasers based on integral high-q resonators in hybrid Si/III-V platforms, Proc. Natl. Acad. Sci. U S A 111 (8) (2014) 2879−2884.
[58] M.L. Davenport, S. Liu, J.E. Bowers, Integrated heterogeneous silicon/III-V mode-locked lasers, Photonics Res. 6 (5) (2018) 468−478.
[59] P.A. Morton, C. Xiang, J.B. Khurgin, C. Morton, M. Tran, J. Peters, et al., Integrated coherent tunable laser (ICTL) with 118 nm tuning range and sub-100 Hz Lorentzian linewidth, in: Proceedings of the Optical Fiber Communications Conference and Exhibition (OFC), IEEE, 2021, pp. 1–3.
[60] C. Xiang, J. Guo, W. Jin, L. Wu, J. Peters, W. Xie, et al., High-performance lasers for fully integrated silicon nitride photonics, Nat. Commun. 12 (1) (2021) 6650.
[61] C. Xiang, W. Jin, J. Guo, C. Williams, A.M. Netherton, L. Chang, et al., Effects of nonlinear loss in high-Q Si ring resonators for narrow-linewidth III-V/Si heterogeneously integrated tunable lasers, Opt. Express 28 (14) (2020) 19926−19936.
[62] S. Liu, X. Wu, D. Jung, J.C. Norman, M.J. Kennedy, H.K. Tsang, et al., High-channel-count 20 GHz passively mode-locked quantum dot laser directly grown on Si with 4.1 Tbit/s transmission capacity, Optica 6 (2) (2019) 128−134.
[63] J.C. Norman, D. Jung, Y. Wan, J.E. Bowers, Perspective: the future of quantum dot photonic integrated circuits, APL Photonics 3 (3) (2018) 030901.

Optical modulators

3

Graham Reed[1], David Thomson[1], Weiwei Zhang[1], Frederic Gardes[1], Lorenzo Mastronardi[1], Ke Li[1], Shinji Matsuo[2], Shigeru Kanazawa[2], Laurent Vivien[3], Christian Lafforgue[3], John E. Bowers[4], Christian Koos[5], Marco Romagnoli[6], Marko Lončar[7], Mian Zhang[7], Stefan Abel[8] and Ling Liao[9]

[1]Optoelectronics Research Centre, University of Southampton, Southampton, United Kingdom, [2]NTT Device Technology Labs, NTT corporation, Atsugi-shi, Kanagawa, Japan, [3]Centre de Nanosciences et de Nanotechnologies, CNRS, Université Paris-Saclay, Palaiseau, France, [4]Department of Electrical and Computer Engineering, University of California, Santa Barbara, CA, United States, [5]Institute of Photonics and Quantum Electronics (IPQ), Karlsruhe Institute of Technology (KIT), Engesserstr. 5, Karlsruhe, Germany, [6]CNIT, Photonics Networks and Technologies Laboratory, Pisa, Italy, [7]John A. Paulson School of Engineering and Applied Sciences, Harvard University, Cambridge, MA, United States, [8]Lumiphase AG, Kilchberg, Switzerland, [9]Photonic Integration, Silicon Photonics Product Division, Intel Corporation, Santa Clara, CA, United States

3.1 Introduction

This chapter discusses the possible mechanisms for optical modulation, technology platforms, the current state of the art, and the prospects for the future progression of optical modulators, all for applications in data communication applications. Due to the ever-increasing demand for data, the transition from electrical to optical interconnect within data centers is inevitable. Due to their inclusion in more traditional optical communications applications, III−V material-based modulators were obvious candidates for short-reach interconnect, but the phenomenal rise of silicon photonics has meant that silicon-based modulators are now extensively deployed within such applications, and a brief review of the emergence of silicon-based modulators serves as a useful introduction to silicon photonics technology. Later in this chapter, we will see that recent alternatives are emerging, where other materials are introduced onto the silicon photonics platform as innovative ways of developing modulators that have the best attributes of the new materials but remain compatible with silicon photonics. The chapter begins with a discussion of modulation mechanisms, key modulator parameters, and figures of merit (FoM).

Integrated Photonics for Data Communication Applications. DOI: https://doi.org/10.1016/B978-0-323-91224-2.00011-4
© 2023 Elsevier Inc. All rights reserved.

3.2 Modulation mechanisms

Optical modulation is obtained by converting an electrical signal into an optical signal. This can be achieved either by directly modulating the current of the light source or by an electro-optic (EO) effect via an optical modulator (Fig. 3.1).

EO effects can be separated into two main categories: *electro-absorption* (EA) and *electro-refraction effects*. The primary electric field effects that are traditionally used in semiconductor materials for causing either EA or electro-refraction are the plasma dispersion effect, the Pockels effect, the Kerr effect, the Franz−Keldysh effect (FKE), and the Quantum-confined stark effect (QCSE). The performance of both EA and electro-refraction optical modulators can be evaluated by defining performance metrics such as extinction ratio (ER), insertion loss (IL), optical bandwidth, footprint, and power consumption. The ER is defined by the ratio between the "on" and "off" states:

$$\mathrm{ER} = 10.\log\left(\frac{I_{\max}}{I_{\min}}\right)$$

where I_{max} is the intensity of the on state and I_{min} is the intensity of the off state (Fig. 3.1).

A large ER is required for long-distance communication, low bit error rates, and high receiver sensitivity. The IL is the optical loss induced by the optical modulator and is defined by:

$$\mathrm{IL} = 10.\log\left(\frac{I_{\max}}{I_0}\right)$$

with I_0 being the optical intensity at the modulator input (Fig. 3.1).

The mechanisms responsible for power consumption in a diode-based optical modulator are the charging and discharging of the device capacitance (dynamic dissipation), the "static" dissipation from the absorption of optical power, and the dissipation from photocurrent flow.

In addition to the ER and IL, the −3 dB electro-optical bandwidth (BW_{3dB}), the maximum data rate (bit/s), device compactness and the optical bandwidth

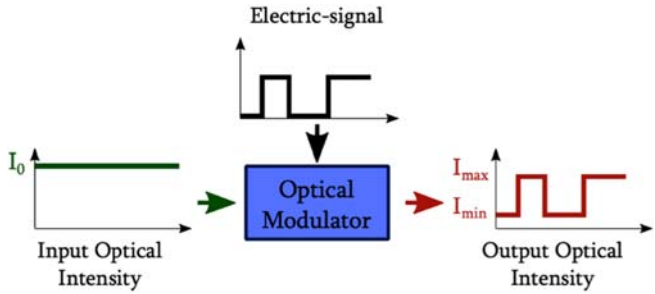

Figure 3.1 Schematic diagram of the principle of an optical modulator, converting an electrical signal into an optical signal.

(or spectral width) referring to the useful operational wavelength range, have also to be determined. More specifically for phase optical modulators (electro-refraction), two additional FoM can be defined: the modulation efficiency ($V_\pi L_\pi$) and the IL for a length, L_π (IL_π). V_π and L_π are the voltage and the length for a pi radian phase change. L_π can be calculated as:

$$L_\pi = \frac{\lambda}{2.\Delta n_{eff}(V_\pi)}$$

where λ is the operating wavelength and Δn_{eff} is the effective index change produced in the modulator under an applied voltage.

The efficiency, $V_\pi L_\pi$, in units of $V \cdot cm$ relates the voltage and length required to obtain a π radian phase shift as defined by:

$$V_\pi L_\pi = \frac{V_\pi \lambda}{2.\Delta n_{eff}(V_\pi)}$$

The lower $V_\pi L_\pi$, the more efficient the modulator.

We have seen that modulation mechanisms are based on the change of refractive index of a material. The most common physical effects associated with electro-refraction modulation are the thermo-optic effect, the plasma dispersion effect, the Pockels effect, and the Kerr effect, while for EA modulation they are the FKE and the QCSE [1].

3.2.1 The thermo-optic effect

The thermo-optic effect is the modification of the real part of the refractive index of the material due to a change in its temperature. This change can be obtained by a Joule effect induced by an electric current. In silicon, this effect is [2]:

$$\Delta n/\Delta T = 1.86 \times 10^{-4}\ K^{-1}$$

Even though this effect is often strong, it is well known to be slow and may require high electrical power to generate heating. In general, it is used as a means to tune the operating point of photonic devices, for example, in Mach−Zehnder interferometers (MZIs) [3] and resonators.

3.2.2 The free-carrier plasma dispersion effect

The free-carrier plasma dispersion (FCPD) effect describes a change in free-carrier concentration in a semiconductor (typically a change in both electrons and holes), leading to a change in refractive index (both real and imaginary). An increase of carriers increases the imaginary refractive index ($\Delta\alpha$: variation of the absorption

coefficient) and decreases the real refractive index (Δn). This effect is described by the Drude−Lorentz model [4]:

$$\Delta \mathrm{n} = -\frac{e^2\lambda^2}{8\pi^2 c^2 \varepsilon_0 \mathrm{n}}\left(\frac{\Delta N_e}{m_e} + \frac{\Delta N_h}{m_h}\right)$$

$$\Delta\alpha = \frac{e^3\lambda^2}{4\pi^2 c^3 \varepsilon_0 \mathrm{n}}\left(\frac{\Delta N_e}{m_e^2\mu_e} + \frac{\Delta N_h}{m_h^2\mu_h}\right)$$

where e is the charge of the electron; c is the speed of light; λ is the wavelength; n is the refractive index of the unperturbed material; ε_0 is the vacuum permittivity; ΔN_e and ΔN_h are the free-carrier concentrations of electrons and holes respectively; m_e and m_h are the electron and hole relative masses; μ_e and μ_h are the mobilities of electron and holes, respectively. For silicon, Soref and Bennet [5] defined a set of empirical equations from experimental results for the wavelengths of 1.3 and 1.55 μm. For $\lambda = 1.55$ μm, the variations of absorption and refractive index with carrier concentration are shown in Fig. 3.2 and take the following form:

$$\Delta n = -8.8\times10^{-22}\Delta N_e - 8.5\times10^{-18}(\Delta N_h)^{0.8}$$
$$\Delta\alpha = 8.5\times10^{-18}\Delta N_e + 6\times10^{-18}\Delta N_h$$

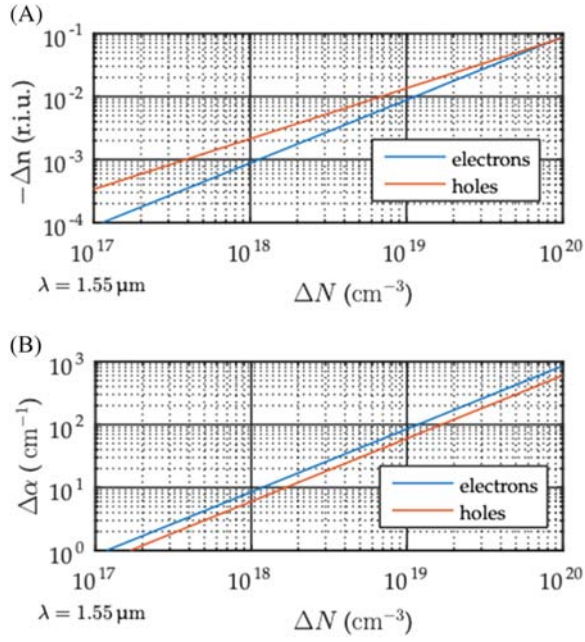

Figure 3.2 Refractive index absorption changes with variation of electron and hole concentrations according to Soref's empirical model [5].

where N_e and N_h are expressed in cm^{-3}, $\Delta\alpha$ in units of cm^{-1}, and Δn in refractive index units.

It can be seen that the variation of hole concentration induces a stronger refractive index change and a weaker absorption change than the variation in electron concentration for carrier concentrations $\Delta N < 8 \times 10^{19}\,cm^{-3}$. The design of Si modulators based on the FCPD effect should take into account these differences to optimize the modulation efficiency and reduce the carrier-induced loss. The modification of electron and hole concentrations in Si is typically carried out by applying a voltage through a diode fabricated in the vicinity of an optical waveguide.

It can also be seen that the EA induced by the FCPD effect is weak [6]. Therefore the FCPD effect is mostly used as an electro-refraction effect, that is, a phase variation. The change in absorption is then limited by design, to minimize loss. The FCPD effect can be achieved through three carrier regimes: injection, accumulation, or depletion of carriers.

- Carrier injection is usually achieved via forward bias of a p-i-n diode leading to current flow through the junction. The intrinsic region is then flooded with carriers (electrons and holes), which induce a refractive index change.
- Carrier depletion is usually exploited via reverse bias of a pn junction. This pn junction induces a small depletion (intrinsic) zone in the vicinity of the junction, that is, a space charge zone. The size of the depletion region increases under reverse bias, sweeping carriers out (depletion), and inducing a variation of refractive index and absorption.
- Carrier accumulation is obtained by employing a MOS-like capacitor. In the Si photonics platform, it is typically composed of a p-doped Si region, an oxide barrier, and an n-doped Si region. When the MOS capacitor is forward biased, charge accumulates at both interfaces of the oxide, holes in the p-doped region and electrons in the n-doped region, inducing a refractive index change close to these interfaces.

FCPD-based optical modulators are fabricated by inserting such electronic structures into Si waveguides. To ensure optimized modulation efficiency, that is, low $V_\pi L_\pi$, waveguide engineering is required to achieve good overlap between the optical mode and the carrier variations. To ensure low IL, care is taken to avoid mode interaction with both doped areas and metallic contacts, and by optimizing the concentrations of the doped regions. The efficiency is linked to the position of all doping regions relative to the waveguide.

Based on the variety of electronic structures to exploit the FCPD effect and the possibility offered by silicon photonics platforms to confine the optical mode, a variety of optical modulators can be then imagined as shown in Fig. 3.3: lateral [7], vertical [8], interleaved [9] structures for carrier injection [10], depletion [7], and accumulation [11]. Even though almost all of them have been experimentally demonstrated, it is obvious that some are more complex to fabricate and integrate with a global circuit, including lasers, modulators, and detectors. However, each has its unique advantages.

3.2.3 Electro-absorption modulation

3.2.3.1 The Franz–Keldysh effect

The FKE occurs in bulk semiconductors. In general, semiconductors absorb photons with energies ($E = hc/\lambda$) higher than the bandgap, shown schematically in

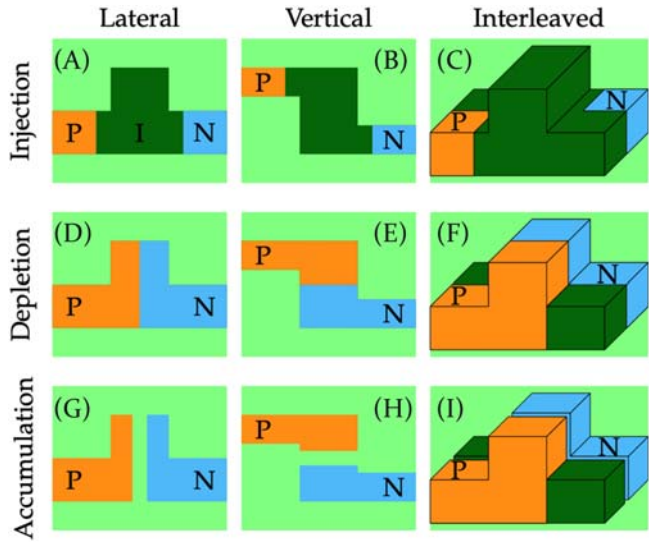

Figure 3.3 Waveguide implementation of modulators based on free-carrier plasma dispersion: (A) lateral carrier injection, (B) vertical carrier injection, (C) interdigitated carrier injection, (D) lateral carrier depletion, (E) vertical carrier depletion, (F) interdigitated carrier depletion, (G) lateral carrier accumulation, (H) vertical carrier accumulation, and (I) interdigitated implementation of carrier accumulation [12,13].

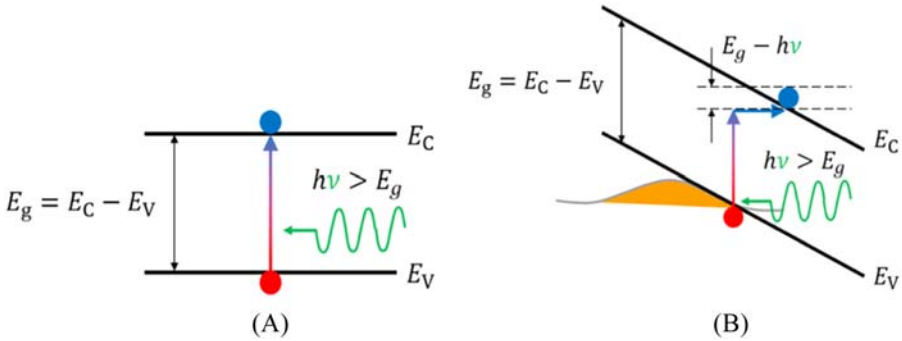

Figure 3.4 Franz−Keldysh effect: (A) photon absorption in equilibrium state, (B) photon absorption under an applied electric field.

Fig. 3.4A. The semiconductor is transparent for photon energies lower than the gap energy (E_g), while it absorbs light for energy higher than the band gap, for example, under an electric field, the energy bands are tilted leading to a penetration of electron and hole wavefunctions into the bandgap, respectively, in the valence and conduction bands, reducing the energy required for a photon to be absorbed by tunneling (Fig. 3.4B). This phenomenon induces a redshift of the absorption spectrum, giving rise to a strong EA effect at photon energies near the gap energy (E_g).

The higher the applied electric field, the greater the electron and hole wavefunction penetration into the energy bandgap, and consequently, the lower the photon energy inducing absorption, and the stronger redshifting. It can also be seen that an electro-refraction effect linked to the EA effect is simultaneously present and is linked via the Kramers−Kronig relations. In silicon the EA due to the FKE (a maximum at $\lambda \approx 1.07\ \mu m$) is weak due to the need for strong phonon-assisted absorption, which is related to the large separation of the indirect and direct bandgap energies. However, in germanium (Ge), this separation is rather low leading to promising characteristics in terms of frequency response, device compactness, and efficiency at a wavelength of about $1.6\ \mu m$, corresponding to the direct bandgap energy. Furthermore, the addition of a fraction of Si into the Ge crystal can successfully shift the operational wavelength to 1550 nm [14].

3.2.3.2 The Quantum-confined stark effect

The mechanism of the QCSE is similar to the FKE, that is, a shift toward lower energy of the absorption spectra (longer wavelengths) under an electric field, but in a confined region. The QCSE is found in multiquantum well (MQW) heterostructures. The confinement of electrons and holes within a well, formed by the combination of two direct-like band-gap semiconductors (one for the wells and another for the barriers) enables a discretization of the density of states above the band edge (Fig. 3.5A). This has two main effects: (1) an increase of the gap energy (shift of the absorption edge to the shorter wavelengths), a sharper absorption edge than in bulk material, and (2) a strong Coulomb interaction between the electrons and holes, leading to excitonic effects. The latter induces an additional absorption peak near the absorption edge. As with the FKE, under an electric field, a redshift of the absorption edge occurs due to the modification of both electron and hole

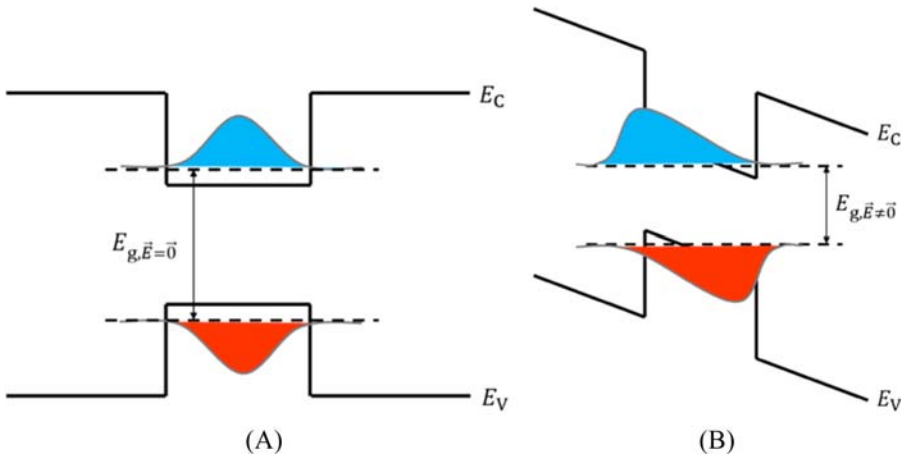

Figure 3.5 Quantum-confined stark effect: quantum well and wavefunction evolutions (A) without an electric field and (B) under an electric field.

wavefunction distributions in the well (Fig. 3.5B) and additionally, a reduction of the excitonic energy contributing to large absorption variation.

Even though III−V semiconductors in QW structures are the material choices due to their direct band-gap nature, Ge-rich SiGe MQW heterostructures have demonstrated promising behavior in terms of efficiency, power consumption, and speed [15,16].

3.2.4 Nonlinear optical effects

Under a high electric field applied to a material, a nonlinear dielectric polarization is created due to the modification of the electronic cloud's center of mass relative to the nucleus. This modification induces a change in refractive index (Δn) as a function of an electric field F:

$$\Delta n \simeq \frac{\chi^{(2)}F}{n} + \frac{3\chi^{(3)}F^2}{2n}$$

where n is the refractive index of the material at the wavelength under consideration. The second-order term related to $\chi^{(2)}$ corresponds to the Pockels effect and consists of the dependence of the refractive index with a static electric field F. The third-order term related to $\chi^{(3)}$ corresponds to the Kerr effect.

3.2.4.1 Pockels effect

The refractive index change induced by the Pockels effect is proportional to electric field. This linear EO effect is largely used in optical communications through ferroelectric materials including lithium niobate (LN) ($LiNbO_3$) [17] and barium titanate ($BaTiO_3$ or BTO) [18]. The absence of a variation of absorption is beneficial for pure phase optical modulation. Such an effect only occurs in noncentrosymmetric materials, and consequently, Si does not exhibit a Pockels effect, unless exposed to material strain that can break its crystal symmetry [19].

3.2.4.2 Kerr effect

The quadratic term in Δn (i.e., the Kerr effect) can also induce nonnegligible refractive index changes in both DC and AC regimes. Indeed, when applying both DC and AC electric fields ($F(t) = F_{DC} + F_{AC} \cdot cos(wt)$), the refractive index change from the Kerr effect can be expressed as:

$$\Delta n = \frac{3\chi^{(3)}}{2n}\left[\left(F_{DC}^2 + \frac{1}{2}F_{AC}^2\right) + 2F_{DC}F_{AC}\cdot\cos(\omega t) + \frac{1}{2}F_{AC}^2\cdot\cos(2\omega t)\right]$$

The term in $F^2{}_{DC}$ is the DC Kerr effect. An additional static refractive index change arises from the AC component with an amplitude proportional to $F^2{}_{AC}$. Furthermore, the AC driving field creates an index variation at double the driving frequency ($2w$) with a quadratic dependence on F_{AC}. Finally, a mixed term

proportional to F_{DC} F_{AC} induces a refractive index change at the same driving field frequency, which is an electric field-induced second-order nonlinear optical effect, that is, a quasi-$\chi^{(2)}$ effect (a Pockels-like effect).

3.2.4.3 Other effects

Any change in the refractive index or absorption of a material under the effect of an external excitation can be used for light modulation. That is the case for many other approaches as follows:

- Electro-refraction and EA based on 2D materials. Graphene is a good example of such approach (see Section 3.6.3).
- Modulation induced by surface wave excitation [20].
- Micro-opto-electro-mechanical systems [21].
- Phase-change materials for low-speed modulation or switching.

3.3 Historical review

The study of silicon photonics modulators goes back almost to the very beginning of research into silicon photonics. The first modulator based on the plasma dispersion effect was proposed by Soref and Bennett in 1986 [22]. It was a theoretical study of a submicron modulator (cross-sectional dimensions), based in silicon, before the popularity of silicon on insulator (SOI) as a substrate material that is the basis of silicon photonics technology today. Modeling suggested that the modulator would achieve a π radian phase shift over a length of 1 mm, for a corresponding loss of 1 dB. This was a p^+-n-n^+ structure that would activate carrier injection under forward bias. The majority of the early silicon photonics work was based on large multimicron waveguides, and consequently, most early modulator designs (and/or switches) were based on carrier injection (e.g., Refs. [23–26]). A detailed review of this very early work is given in Ref. [27]. By the mid-1990s, almost all modulator work focused on designs in SOI material, whereas some of the earlier work had been based on a doped silicon cladding. Authors were also investigating 3-terminal devices (e.g., Refs. [28,29]), or devices with very deep doping (e.g., Refs. [30,31]) to improve modulation speed and power efficiency, or even to look at new applications such as beam steering. Optimization of both 2-terminal and 3-terminal modulators continued throughout the 1990s and early 2000s, to continue to improve device performance in these multimicron waveguide-based devices, with speed, efficiency, and power consumption still being the focus (e.g., Refs. [32–34]). In 2000 Dainesi et al. [35] produced a new device (Fig. 3.6), also targeting performance optimization. While the performance was modest at 10 MHz, the device was interesting as it included segmented doping regions which would return to popularity in later years. The device was fabricated in a large 10-μm thick silicon guiding layer and based on carrier injection, but this was a time when researchers were just beginning to recognize the high speed-potential of high-speed silicon modulators, which could be achieved by reducing the cross-sectional dimensions of

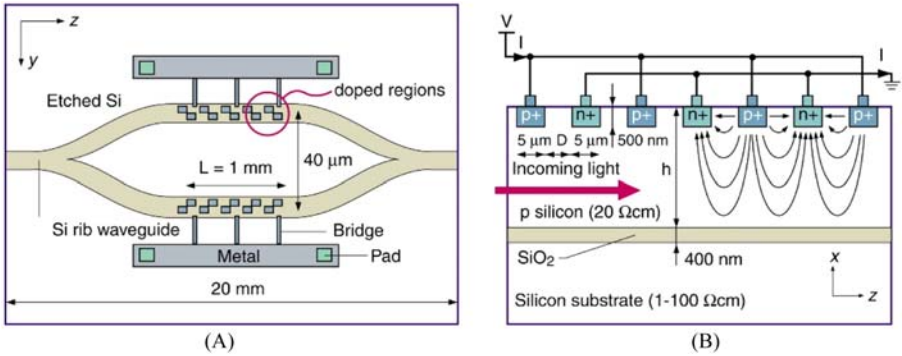

Figure 3.6 Modulator of Dainesi et al., with segmented doping regions from 2000, in (A) plan view, and (B) cross sectional view [35].

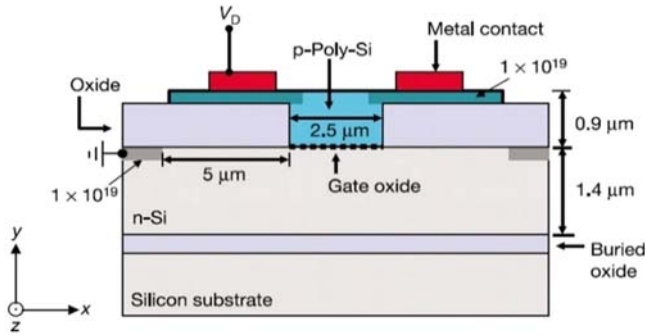

Figure 3.7 The first silicon modulator to operate beyond 1 GHz [35].

the waveguides. Consequently, the typical multimicron devices that were common in the 1980s and 1990s began to give way to smaller devices, toward the submicron devices that dominate today.

The year 2004 was important in the evolution of silicon photonics modulators. In February of 2004, the Intel group published an important breakthrough paper [36], reporting a speed milestone. This was the first experimental device to exceed a bandwidth of 1 GHz, but also the first carrier accumulation-based modulator (also known as a MOSCAP device). The device was demonstrated with a 1 Gb/s data stream. The device is shown in Fig. 3.7 and is notable for high-speed performance despite the overall device thickness exceeding 2 μm. The device was significantly improved the following year [37], to reach data rates of 10 Gb/s at an ER of 3.8 dB.

Also, in 2004, Png et al. [38] designed a range of carrier injection devices predicted to operate from 70 MHz to beyond 1 GHz for the first time. These devices were based on a waveguiding layer of 1.38 μm in thickness. Furthermore, this chapter proposed the use of the preemphasis drive technique to improve the speed to almost 6 GHz. The devices discussed thus far in this introduction have all been phase

modulators that realized intensity modulation (IM) via incorporation into an MZI. However, in 2005, Xu et al. [39] very effectively combined a ring resonator structure with a small waveguide to produce a low-power modulator operating at 1.5 Gb/s. This device, however, while being a carrier injection modulator, was based in a much smaller waveguide, with a width of 450 nm, and a height of 250 nm, much more typical of the small waveguides that are common today. The device was later enhanced in speed to 18 Gb/s in 2007 using the preemphasis technique proposed by Png [38]. The attributes of resonant modulators vs interferometric devices such as the MZI will be discussed in more detail later, but it is worth noting in passing that resonant structures typically offer the possibility of smaller footprint than MZIs, at the expense of narrow spectral range and thermal sensitivity. The latter has limited the commercial applications of ring structures, but devices are now starting to emerge from the research laboratory, typically utilizing an electronic feedback circuit to compensate for thermal drift, adding complexity, additional energy consumption, and potentially a need for additional electronic integration. EA modulators (EAMs), on the other hand, also offer a small footprint, potentially without the need for electronic control, although at the price of fabrication complexity. EAMs are discussed in more detail in Section 3.4.3. The trend to smaller waveguides throughout the 2000s opened up the possibility of using depletion of carriers, rather than injection or accumulation, because it became possible to deplete an entire waveguide if it is small enough. Depletion modulators are the most common commercial device available today, and almost the workhorse of silicon photonics transceivers. The first proposal for a depletion modulator in a submicron waveguide was published in 2005 by Gardes et al. [40], an important milestone in silicon modulator development. The device design was based on a horizontal p-n junction across the waveguide (Fig. 3.8) and was

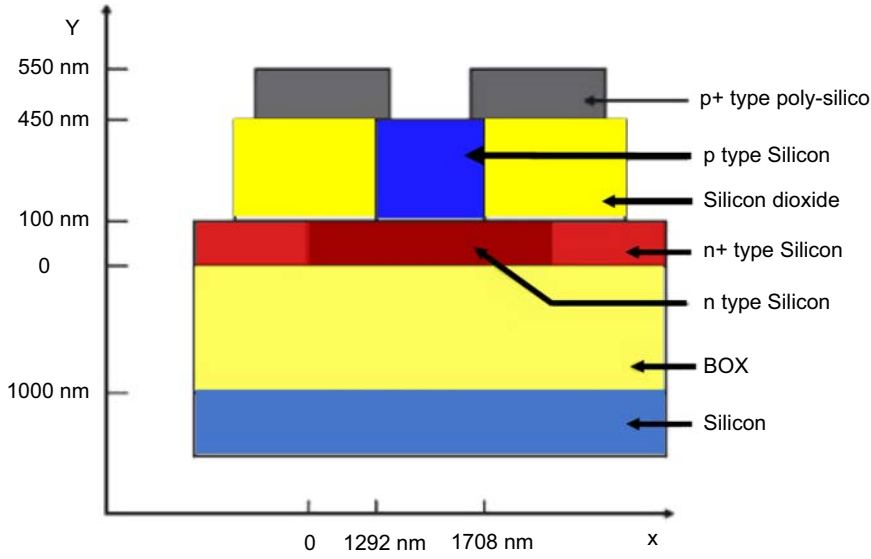

Figure 3.8 The first proposed depletion modulator, now an industry standard device [8].

analyzed via simulation, predicting an unprecedented 50 GHz bandwidth. A series of important results were reported within the next 6 years to take experimentally verified data rates to 30 Gbit/s [41] and 40 Gbit/s [42] in 2007, 40 Gbit/s with high ER in 2011 [43], 40 Gbit/s for transverse electric (TE) and transverse magnetic (TM) polarizations also in 2011 [8], and 50 Gbit/s in 2012 [44].

Since then a plethora of research has been published to increase the data rate of silicon photonics transceivers, while reducing power consumption, integrating with electronic drivers, and utilizing enhanced modulation schemes to improve the overall performance. Today silicon photonics is unrecognizable when compared to the preliminary activities in the late 1980s and early 1990s, and the modulator is perhaps the device that has changed the most, moving from modulation speeds of a few Mbits/s to more than 100 Gbit/s in 2020 [45], at dramatically lower power consumption, and increasing complexity both in terms of integration with other devices, but also in terms of integration with other materials. This chapter discusses these changes, concluding with comments on the outlook for the future.

3.4 Group IV modulators

In the previous section, many of the EO effects used in optical modulators were introduced together with metrics used to assess the performance, such as the ER, IL, 3 dB EO bandwidth, maximum data rate, optical bandwidth, and compactness. The efficiency of the modulator in terms of its voltage and current-length products was also discussed. In this section, more details of the practical implementation of devices based upon these effects using group IV materials are provided, together with demonstrated examples. Another important aspect to consider is the required interfacing electronics, namely, the driver amplifier. As the performance requirements have advanced, it has become increasing important to consider the driver, modulator, and any parasitics from the chosen integration scheme, as a single system to ensure that the highest performance levels are achieved. This is discussed in the last part of this section.

3.4.1 Modulation methods

Depending on the application of the data link that the modulator is serving, either phase or IM may be required. In the previous section, we have seen that the thermo-optic, FCPD, Pockels and Kerr effects are predominantly electro-refractive whereas the FKE and QCSE are predominantly electro-absorptive. Although the FCPD effect is predominantly electro-refractive, it has been the most popular effect used for IM owing largely to the simplicity in forming devices of this type in a silicon waveguide system by using CMOS-like process steps. To achieve IM from electro-refractive effects, a passive optical structure is employed to convert phase modulation. Many structures have been proposed and demonstrated to perform this

function that can be split into two categories, interferometric and resonant, the operation of which will be described next as well as in Chapter 5.

3.4.1.1 Interferometric photonic modulators

The most popular structure using the interferometric approach is the MZI as shown in Fig. 3.9. The operation of interferometric modulators is based upon splitting the input light into two paths, performing phase modulation of the light in at least one of the two paths before recombining them into one or more output paths.

When light from the two paths recombines, it will interfere, depending on the relative phase. If in-phase, the light will interfere completely constructively giving a maximum intensity output. If π radians out of phase, it will interfere completely destructively and give a minimum intensity output. If the relative phase is in between these two states, then the intensity will be between the minimum and maximum, following the well-known transfer function as shown in Fig. 3.9.

In this way the phase modulation produced by the phase modulator can be translated into IM. In a perfect interferometer the light is split equally into the two paths, experiences the same propagation losses in each path, and is then recombined. Deviation from this ideal case will cause the minimum level to deviate from zero and increase loss at the maximum level. Ideally, the device will be insensitive to the polarization and the wavelength of light, which may be possible within certain practical operational ranges. The phase of one or both of the arms may be varied. When the phase in both arms is varied, it can be varied in opposition to effectively double the relative phase difference, at the combiner. Using this approach has many advantages, including lower drive power, shorter device length, reduced output chirp, and balanced optical losses. However, it requires double the number of drive signals unless a single-drive push−pull approach is employed [46].

3.4.1.2 Resonant photonic modulators

The most popular structure using the resonant approach to convert phase shift to amplitude modulation is the ring resonator as shown in Fig. 3.10. Resonant modulators work by incorporating the phase modulator into the resonant structure so as to enable modulation of the optical path length, and therefore the wavelength at which

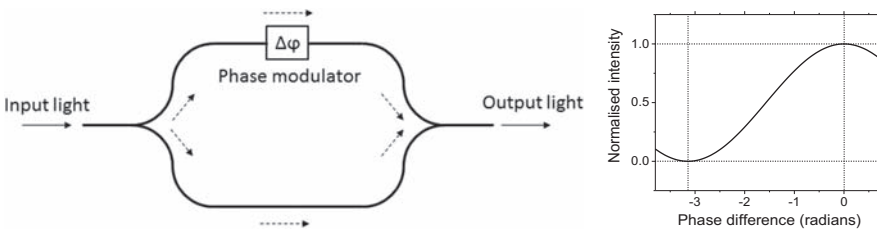

Figure 3.9 Diagram of a Mach−Zehnder interferometer-based modulator and transfer function.

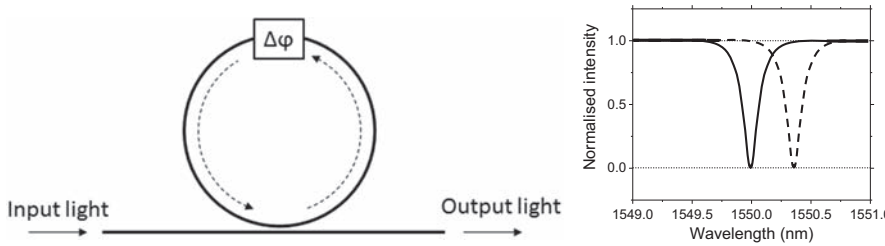

Figure 3.10 Diagram of a ring resonator-based modulator.

the device resonates. When the ring resonator is at resonance, a minimum in transmission results, and away from resonance a maximum occurs. In this way, by inputting light into the device which is initially at resonance and then using the phase modulator to shift the resonant wavelength of the ring away from that of the input light, the transmission can be moved from a minimum to a maximum.

The transfer function of the ring resonator with two different drive conditions is shown in Fig. 3.10. We show the case where the resonance depth is large, which occurs at critical coupling, where the ring round-trip losses are such that the intensity in the ring at the coupling point equals the light coupled into the ring. If these intensities are not balanced, the ring may be under or overcoupled and the depth of the resonance reduced. The quality factor (Q) of the ring is a very important parameter which is dependent on the round-trip losses. A large Q will provide a narrow resonance meaning that only a small shift in the resonance wavelength is required to achieve a large degree of IM. A large Q, however, implies that the photon cavity lifetime will be large restricting the speed of modulation. Therefore a compromise of the Q is required to have both efficient modulation and high-speed operation.

3.4.1.3 Interferometric versus resonant approaches

Table. 3.1 summarizes the comparisons between MZI and ring modulators. In comparison with MZI modulators, the ring modulators benefit from compact size, resulting in a lower capacitance that can be driven as a lumped element with higher bandwidth, lower power consumption around 10 seconds of fJ/bit. However, resonant devices are wavelength dependent, more sensitive to both fabrication errors and temperature changes implying that more active tuning would be required to hold the device at the correct operating point. The power required for the control circuitry and tuning mechanism used to stabilize will therefore likely be much larger. The overall power consumption will depend strongly on how the device is to be operated and the environment in which it is required to operate.

3.4.2 Examples of interferometric and resonant devices employing the Free-carrier plasma dispersion effect

In this section, examples of interferometric and resonant devices based on carrier injection, depletion, and carrier accumulation are presented.

Table 3.1 State of the art of different types of silicon modulators.

	Mach–Zehnder interferometer modulators	**Resonant photonic modulators**
Compactness	> 1 mm	<20 μm*2π
Wavelength dependence	Broadband	*Discrete resonances, <0.5 nm*
Temperature sensitivity	Less sensitive	Highly sensitive
Insertion loss	High	*Low*
Sensitivity to fabrication error	Less sensitive	Higher sensitivity
Power consumption	>1 pJ/bit	*10 s fJ/bit*
Capacitance	High	*Low*
Bandwidth	Low	*High*

The overall performance of a given device should be assessed taking into account the different performance metrics, efficiency, IL, operating speed, optical bandwidth, etc. Many of these metrics can be traded off against one another. For example, increasing the carrier doping levels generally increases the operating speed, but also increases optical loss. A figure of merit taking into account all of the considered performance metrics can be used to give an overall picture; however, this may not necessarily give an indication of the best device for a given application. Therefore we summarize the current state of the art in Table 3.2, for some of the critical performance metrics to give an idea of the bounds of performance achievable.

The state of the art of carrier injection modulators allows data rates up to 90 GBaud but with external circuits for equalization to extend the bandwidth, indicating that the device performance is not comparable with carrier depletion modulators in terms of efficiency and bandwidth, as summarized in Table 3.2. The carrier depletion and carrier accumulation approaches have been widely studied as they are the practical options for high-speed data modulation without any (or with less complicated) preemphasis drive techniques. The carrier accumulation modulators were first reported with 1 Gb/s data rate in 2004 and have experienced a large improvement to reach 40 Gb/s in 2014 by using horizontal MOS junctions with CMOS drivers and 60 Gb/s by using vertical MOS junctions by external driver probing method. The experimentally demonstrated vertical MOS junction modulators are compatible with popular thin SOI photonics platforms (with mature silicon depletion modulators) and show promise to reach high-speed, compact, lumped MZI modulators in comparison with depletion MZI modulators typically with traveling wave electrodes. Silicon depletion modulators are the most widely studied devices that have been optimized with data rates beyond 100 GBaud [56] and co-integrated with CMOS drivers [45].

3.4.3 Electro-absorption modulators

The major EA modulation mechanisms such as the FKE in bulk material and QCSE in QW layers have been demonstrated in bulk germanium and GeSi alloys. EA

Table 3.2 State of the art of different types silicon modulators.

References	Year	Optical structure	Doping structure	Efficiency (V · cm)	Insertion loss (dB)	ER (dB)	V_{pp} (V)	Bitrate (Gb/s)	Size
Carrier injection									
Xu et al. [10][a]	2007	Ring	PIN	–		9	8	12.5	5 μm *2π
Meister et al. [47]	2013	Nano cavity	PIN	0.014	6	10	2.7	12.5	8.8 μm
Akiyama et al. [48][a]	2013	MZI	PIN	0.01	7.2	4.3	4.35	50	250 μm
Shakoor et al. [49]	2014	Nano cavity	PIN	–	1	8	0.5	3	3.5 μm
Sobu et al. [50][a]	2021	MZI	PIN	2	2.78	3.1−1	7.2−2	50−90	1000 μm
Carrier accumulation									
Liu et al. [36]	2004	MZI	Horizontal MOS	8	6.7	16	7.7	1	2.5 mm
Van Campenhout et al. [51]	2012	Ring	Horizontal MOS	0.24	3	8	1.5	3	2.5 μm *2π
Webster et al. [52]	2014	MZI	Horizontal MOS	0.2	–	8	1	40	400 μm
Sodagar et al. [53]	2014	Microdisk	Horizontal MOS	–	5	8	4	15	3 μm *2π
Debnath et al. [54]	2018	MZI	Vertical MOS	1.53	2.5	3.65	6	25	500 μm
Zhang et al. [55]	2021	MZI	Vertical MOS	1.8	–	–	4	60	200 μm
Carrier depletion									
Liao et al. [42]	2007	MZI	Horizontal PN	4	4	1.1	6	40	–
Thomson et al. [44]	2012	MZI	Vertical PN	2.8	3.7	5.5	6.5	50	1 mm
Li et al. [45][b]	2020	MZI	Vertical PN	1.5	6.9	3	∼3	>100	2.47 mm
Sakib et al. [56]	2021	Ring	L-shape PN	0.53	9	3.9	0.8	128	6 μm *2π

MZI, Mach−Zehnder interferometer.
[a]Electrical driving signals are with preemphasis.
[b]CMOS driver codesigned for high bandwidth silicon modulator.

modulation is obtained for a narrow wavelength range by varying the imaginary part of the refractive index with an electric field as theorized independently by Franz [57] and Keldysh [58] and, therefore, the absorption coefficient of the material becomes proportional to the electric field strength. Thanks to the germanium direct bandgap around 0.8 eV (corresponding to 1550 nm), FK-based transceivers operating in the C-band telecommunication window can be integrated to a silicon photonic platform. GeSi with a low Si content ($\sim$1.5% [14]) is often used to tailor the operation point by varying the direct bandgap energy and/or by engineering the strain accumulated, when the alloy is grown [59]. More recently, the engineering of epitaxial growth techniques allowed deposition of QWs to exploit the QCSE in waveguides utilizing Ge, GeSi [16,60–65]. QWs allow for a wider wavelength range of operation, a graphical comparison of the energy bandgap for bulk material and QW is depicted in Fig. 3.11. In a bulk material the energy bandgap can only be varied by changing the material composition and/or applying strain. In a QW the electronic behavior can be tailored and the nanostructure thickness affects the bandgap energy as illustrated in Fig. 3.11 (arrows from orange to blue).

The benefit of using QCSE-type devices becomes clear as deposition engineering allows shifting of the operation band of such a modulator from the O-band to the C-band by modifying strain and quantum confinement [67].

The integration of EAMs into a CMOS photonic platform is enabled through the formation of a PIN diode (Fig. 3.12) integrated into a portion ($<$100 μm) of the waveguide. The diode is driven in reverse bias as a lumped element circuit. When no bias is applied, the propagating optical wave is partially absorbed due to the built-in electric field, whereas at reverse bias the absorption increases, where the

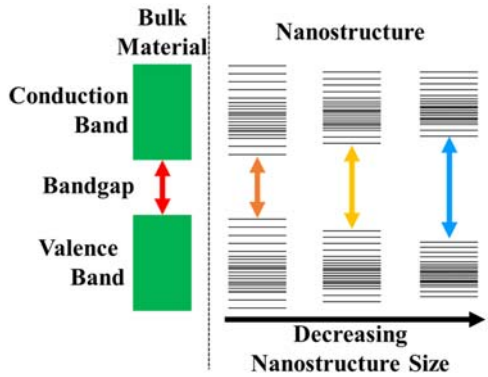

Figure 3.11 Comparison of the bandgap energy for bulk and QW semiconductors. In bulk the bandgap is fixed and depends on the material composition, it can be varied with external forces such as strain. In QWs the quantum confinement increases the bandgap if the size of the nanostructure reduces [66]. *QW*, Quantum well.
Source: Reproduced from L. Mastronardi, GeSi Franz-Keldysh modulator for silicon photonic integrated circuits (Doctoral thesis), University of Southampton, 2019. <https://scholar.google.co.uk/scholar?hl = en&as_sdt = 0%2C5&q = GeSi + Franz-Keldysh + modulator + for + silicon + photonic + integrated + circuits&btnG = > (accessed 09.10.21).

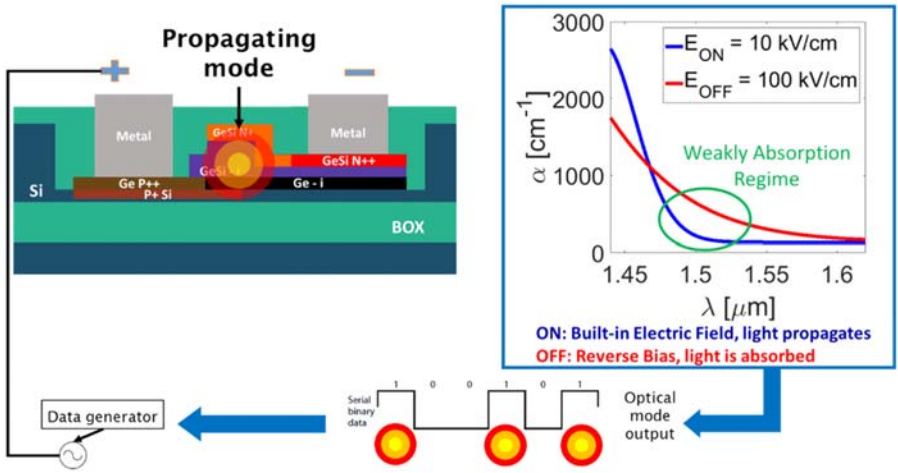

Figure 3.12 PIN FKE modulator cross section and absorption coefficient versus wavelength with (red) and without (green) reverse bias. The PIN diode is integrated to a waveguide, where the mode propagates through the intrinsic region of the PIN junction. *FKE*, Franz−Keldysh effect.

difference in absorption is proportional to the strength of the electric field and length of the modulator.

Typically, EAM FKE-based devices have a relatively small footprint in the region of 60 μm^2, resulting in a low power consumption (<50 fJ/bit) and display data rates in excess of 50 Gbps with the possibility to operate in TE and TM polarizations. The operational bandwidth is typically in the order of 30 nm within the C-band. QCSE modulators [60] are also implemented using PIN diodes, where the stronger field effect and the extended operation regime are counterbalanced by a more complex integration process compared to FKE modulators.

The difficulties of integration of QCSE structures are linked to the mode matching and interfacing of a relatively large QCSE QW stack waveguides to a smaller silicon waveguide. A recent demonstration of integration of QSCE QW in silicon waveguides is illustrated in Fig. 3.13.

Table 3.3, summarizes a range of integrated EAMs using Ge [68,69] and operating in the L-band as well as EAMs using GeSi [14,70,71] with varying operational points in the C-band. Demonstrated EAMs have reached data rates beyond 56 Gb/s with low power consumption (12.8 fJ/s), low dark current (< 47 nA), and compact design (down to 16 μm^2 [68]). The optical bandwidth ranges between 14 [18] and 35 nm [68,69,71]. Coarse wavelength division multiplexing (CWDM) [72] increases the optical bandwidth up to 80 nm at the cost of speed, limiting the speed in this example to 10 Gb/s, and power consumption (5.82 mW per channel).

In Table 3.4, QCSE-based devices are summarized that have been demonstrated to operate in the E-band [16,61] or C-band [62,63] with low power consumption (8.25 fJ/bit [63]) and compact footprint (120 μm^2 [7]). Recently, waveguide-integrated

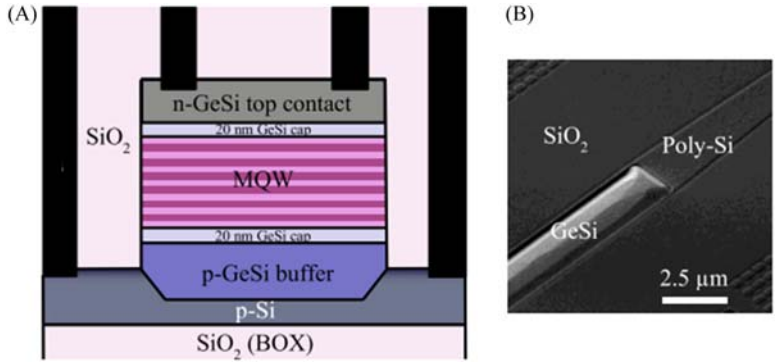

Figure 3.13 (A) Cross section of the QCSE EAM. (B) SEM image of the waveguide QCSE EAM coupled to a poly-Si waveguide. *EAMs*, Electro-absorption modulators; *QCSE*, Quantum-confined stark effect; *SEM*, Scanning electron microscope.
Source: Reproduced from C. Porret et al., "O-band GeSi quantum-confined stark effect electro-absorption modulator integrated in a 220nm silicon photonics platform," in: *2020 IEEE Symposium on VLSI Technology*, 2020, pp. 1−2, https://doi.org/10.1109/VLSITechnology18217.2020.9265082.

QCSE EAMs operating in the O-band [60] and C-band [64,65] have been proposed, showing wide optical operation range, low power consumption, compact design, and a relatively large ER.

3.4.4 Modulator driver amplifiers

In the optical transceiver link, a modulator driver operates as a broadband amplifier and sees the optical modulator as a low-impedance load. The driver amplifier accepts a low-voltage swing signal from the preceding stage, such as a multiplexer or predriver, and amplifies the signal up to several volts. There are many metrics when evaluating the performance of a driver amplifier, and the most common specifications are bandwidth and group delay, voltage swing, and power efficiency.

3.4.4.1 Bandwidth and group delay

The bandwidth of the typical broadband driver can be defined as the 3-dB bandwidth of the voltage gain (S21 in dB). Meanwhile, the group delay, which is a measure of phase distortion over the frequency range of interest, plays a dominant role when the bandwidth of the amplifier extends into the millimeter wave range. Most driver amplifiers need to cope with broadband signals ranging from several kHz to tens of GHz. Owing to the adoption of on-chip peaking inductors/transformers, it is possible to broaden the gain-bandwidth product of the amplifier. However, it is much more challenging to maintain a constant group delay over the entire frequency range, especially when the modulator is integrated. An example of this

Table 3.3 Recent progress in silicon Franz−Keldysh modulator.

References	λ (nm)	Area (μm^2)	3db band (GHz)	Speed (Gb/s)	ER dynamic (dB)	Voltage (V)	Dark current (nA)	Energy (fJ/bit)	IL (dB)	ER (dB)
Liu et al. [70]	1550	0.6×50	1.2	–	–	3 (V_{pp}) − 4 (bias)	< 65	50	7.5	10
Feng et al. [71]	1550	0.8×50	38	28	4.5	3 (V_{pp}) − 3 (bias)	15.4×10^3	147	4.8	5.9
Lim et al. [68]	1600	16	–	1.25	–	−5 (bias)	–	–	9.6	10
Srinivasan et al. [69]	1610	0.6×40	>50	56	3.29	2 (V_{pp}) − 2 (bias)	<47	12.8	4.9	4.6
Mastronardi et al. [14]	1566	1.5×40	56	56.2	5.2	4 (V_{pp}) − 2.7 (bias)	2×10^3	44	10.6	5.7

ER, Extinction ratio.

Table 3.4 Recent progress in silicon Quantum-confined stark effect modulator.

References	λ (nm)	Area (μm^2)	3db band (GHz)	Speed (Gb/s)	ER dynamic (dB)	Voltage (V)	Energy (fJ/bit)	IL (dB)	ER (dB)
Rouifed et al. [16]	1400	3×50	–	–	–	4 (V_{pp}) −4 (bias)	22	3	4
Chaisakul et al. [61]	1433 1444	3×90	23	–	–	1 (V_{pp}) −3, −4 (bias)	108	5.5 12	10
Edwards et al. [62]	1550 1490	0.6×200	–	–	–	1 (V_{pp})	–	3.5−4.5 6.5−8	6.5 15
Roth et al. [63]	1541	225×625	–	–	–	1 (V_{pp}) −1 (bias)	8.25	–	3
Ren et al. [64]	1425 1548	8	–	7 @ 1460 nm (oscilloscope trace)	3.2 @ 1460 nm (oscilloscope trace)	1 (V_{pp}) −6.5 (bias)	–	–	2.84 (photocurrent)
Zang et al. [65]	1450 1550	100×2	–	–	–	1 (V_{pp}) −1, −3 (bias)	–	–	∼1.4 (photocurrent)
Porret et al. [60]	1350	$\sim 2 \times 40$	–	–	–	1 (V_{pp}) −1 (bias)	–	16	8

ER, Extinction ratio.

dilemma is demonstrated in Ref. [73] where an outstanding bandwidth from DC to 108 GHz had been delivered, whereas its group delay variation is up to 80 ps over that frequency range. Therefore we have seen that most commercial driver amplifiers can only deliver a usable bandwidth up to 70 GHz, which is approximately equal to a symbol rate at 100−120 GBaud.

3.4.4.2 Voltage swing and power efficiency

The voltage swing of a driver is normally expressed as "V_{pp}" for a single-ended output voltage swing or "V_{ppd}" for differential outputs. Traditionally, these figures are measured by feeding the output node of the driver into a digital oscilloscope, which is assumed to be exactly 50 Ω. Commonly, a driver should provide a single-ended output swing up to several volts [74−76]. The power consumption (in mW) divided by the operation speed (Gb/s) generates a figure for power efficiency (mW/Gb = pJ/bit). Obviously, a higher voltage swing drives more current and hence comes with an inferior power efficiency (energy/bit). This figure of merit is typically used to evaluate the performance of a driver. However, these figures become meaningless when the driver is custom designed for a bespoke modulator. For example, the drivers in Ref. [77] designed for ring resonators see the load as a lumped capacitor, whereas some drivers are designed for traveling wave modulators and see a load significantly less than 50 Ω [78]. Furthermore, the adoption of complex modulation formats makes it even harder to make a fair comparison, as complex modulation formats always come with a lower signal-to-noise ratio. Therefore it has become common to combine the optical ER and power efficiency figure: power efficiency = power/(data rate*ER) [79]. This figure provides a relatively fair comparison of different types of modulator and different modulation schemes.

3.4.4.3 Circuit topology

When it comes to practical implementation, the circuit topologies of the driver amplifier could be classified into two categories: voltage [80−83] and current modes [75,76]. In both approaches, stacked transistors are adopted to protect the switching transistors from voltage breakdown issues and to withstand the high voltage swing.

Fig. 3.14A shows that the voltage-mode approach takes advantage of the CMOS process by configuring the PMOS and NMOS transistors as push−pull amplifiers. During the operation, either the PMOS or NMOS is switched on at any one time hence eliminating static current flow. Therefore this configuration is preferred due to its potential for lower power consumption. However, the switching behavior of PMOS and NMOS leads to a dynamic change of the output impedance when transistors operate in different operating regions. Even with a built-in impedance control module [81], it is hard to maintain a constant output impedance over a broadband frequency range. As a result, the voltage-mode approach is usually utilized to drive lumped elements, such as ring resonators [80,82] or small segments within a Mach−Zehnder modulator (MZM) [79,84]. Furthermore, another unavoidable issue

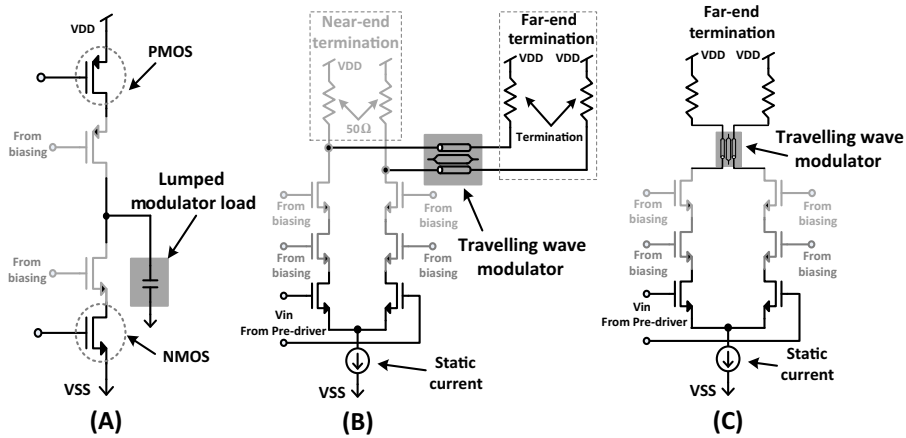

Figure 3.14 (A) Voltage-mode push−pull driver integrated with lumped element. (B) Current-mode driver amplifier with double termination. (C) Open-drain driver amplifier integrated with traveling wave modulator.

is the relatively low transition frequency (*ft*) of the PMOS transistor, which becomes a major bottleneck for higher speed applications. The highest baud rate of reported voltage mode drivers can be found in Ref. [83], which is realized on the 28 nm FD-SOI process and operates up to 60 GBaud [on−off keying (OOK)].

In contrast to the push−pull amplifier, as shown in Fig. 3.14B, the current-mode approach eliminates the PMOS transistors and hence avoids the speed limitation issue. By switching the NMOS transistors within a CMOS process or NPN bipolar transistor in a BiCMOS process, the current is switched between two pairs of termination resistors (near-end 50-Ω resistors and far-end termination resistors) and hence a reliable impedance matching over the broadband frequency range is provided. Therefore current mode has become the primary choice for modulators that are built in a traveling wave configuration (such as the MZM). The major concern with the current mode approach is the relatively low power efficiency, which comes as a result of needing to deliver current flow to four parallel load resistors.

Therefore, as shown in Fig. 3.14C, a popular solution is to eliminate the two near-end termination resistors and hence significantly reduce the power consumption. This kind of configuration is called an "open-drain amplifier" and has been adopted in many recent demonstrations [78,85,86]. However, the precondition is the precise control over device packaging so that the parasitic effects (from wire bonding and/or flip-chip bonding) that induce impedance discontinuities are minimized. Even if the impedance discontinuity issue can be alleviated, a major concern is that the system bandwidth is predominantly determined by the bandwidth of the electrodes of the traveling wave modulator. An example is reported in [78], in which two segments of traveling wave modulator are monolithically integrated with two separate open-drain amplifiers. The length of the first segment is twice that of the

second, and the first segment E-O bandwidth is significantly lower (by more than 50%) than the second one.

3.4.4.4 Codesign of driver and modulator

To tackle the demanding technical challenges in the field, the modulator and driver amplifier must be considered a single integrated system. This means that neither the amplifier nor the optical modulator can be treated as a standalone component, and the realization of system functionality depends on the integration of these functions at the design stage. Next, we highlight the three most common schemes that have been developed in the field.

1. Segmented modulator integrated with driver amplifier

Segmented modulators, codesigned and copackaged with a number of driver amplifiers, offer a variety of advantages for high-performance transmitters, where advanced modulation formats [22,36–40] (such as PAM-N and QAM-N), high ER [79,87], or optical signal shaping (such as optical feed-forward equalization) [88] have become the system requirement. As shown in Fig. 3.15A, the phase shifter is divided into a number of segments, with each segment fed with an independent electrical signal. With such a configuration, each driver unit only sees a small portion of the phase shifter, which is seen as a lumped element to the driver amplifier, and hence the power-hungry termination resistors can be eliminated.

However, a major challenge for this scheme is that the timing of the electrical signals needs to be matched with the propagation delay of the optical signal. As a result, a timing control function (τ) is always built within each electrical driver amplifier. In practical implementations, this time delay function could be realized with a passive device, such as an on-chip transmission line [87], or with active components, such as a tunable delay element [89]. Both schemes introduce additional overhead (power and chip area) to the system and suffer from the unavoidable process–voltage–temperature variations.

2. Inductive peaking

Inductive peaking is a standard technique that has been utilized for bandwidth enhancement in numerous high-speed electrical circuit designs. When a driver amplifier sees a photonic device configured as a lumped element, the device capacitance creates a low

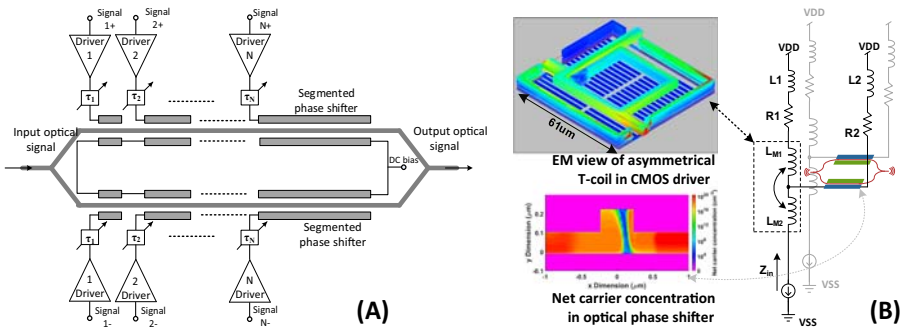

Figure 3.15 (A) Conceptual model of segmented modulator integrated with driver amplifier. (B) T-coil peaking embedded in between the driver and traveling wave modulator.

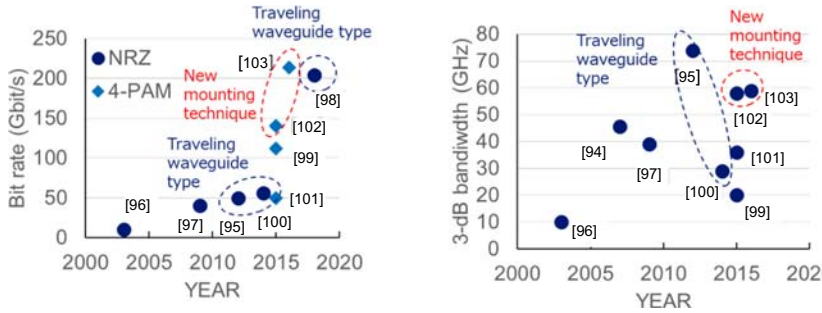

Figure 3.16 Bitrate and 3-dB bandwidth change of EA modulators [91–100]. *EA*, Electro-absorption.

impedance load at high frequency. When a driver amplifier is integrated with a traveling wave modulator, the radio frequency (RF) loss of electrodes could significantly degrade the system bandwidth. In either case, a precisely modeled peaking inductor between the photonics and electronic devices can dramatically boost the system bandwidth. This inductor could be realized with a bonding wire [80] or with T-coil peaking [45,90]. As shown in Fig. 3.16B, an on-chip asymmetrical T-coil is embedded between the driver amplifier and the traveling wave modulator. According to Ref. [45], such a configuration can offer an additional 70% bandwidth enhancement over traditional peaking inductors and also demonstrates that copackaged driver and modulator could operate up to 112 GBaud. The concerns with this technique are the additional chip area on the driver chip and the variation of group delay at high frequencies.

3. Nonlinearity cancellation

Due to the raised cosine transfer function of the MZI, nonlinearity issues always exist when complex modulation formats (PAM-N/QAM-N) are utilized. To address this concern, the segmented MZI design [101] has introduced additional segments and has been integrated with a corresponding driver to compensate the nonlinearity. In Ref. [81] a driver that is integrated with a traveling wave modulator has incorporated an additional predistortion function to compensate the nonlinearity. The nonlinearity issue also exists within ring resonator modulators. The Lorentzian shape of E-O transfer characteristics leads to a wavelength-dependent static nonlinearity. Different bias voltages also dynamically change the small signal bandwidth which produces a dynamic nonlinearity. Recent work from Intel [77,80] has introduced nonlinear feed-forward equalization within the driver circuit itself and through this approach, an experimental demonstration of a 112-Gbps PAM-4 signal, with a measured Transmitter Dispersion Eye Closure Quaternary (TDECQ) less than 1 dB has been possible.

3.5 III–V modulators

In this section, we describe III–V compound semiconductors modulators for high-speed optical communication systems. A directly modulated laser (DML) is the simplest method of generating the signals by changing the injection current,

providing a low-cost and low-power solution. However, there are limitations such as modulation speed and frequency chirp (i.e., transmission distance). EA and MZM III−V modulators require separate light sources but overcome these drawbacks. Table 3.5 compares vertical-cavity surface-emitting lasers (VCSELs), distributed feedback (DFB) lasers, EA modulators, and MZMs. Directly modulated multimode VCSELs are typically limited to less than 50 Gbit/s and the transmission distance is limited to ∼100 m. DFB lasers provide greater power and extend the transmission range, although they are limited by frequency chirp when directly modulated. III−V modulators use an MQW, where the detuning between the operating wavelength and the exciton peak wavelength is ∼50 nm for the EA modulator and ∼150 nm for the MZM modulator. The EA modulator with a several-hundred-μm length is suitable for IM with a bitrate of over 50 Gbit/s by using the QCSE. As the detuning increases, it can be used as a phase modulator with small IM. Thus III−V MZMs are suitable for generating signals with advanced modulation formats such as quadrature amplitude modulation (QAM) that can increase throughput and transmission distance. However, because the modulation efficiency is low and the length of the phase modulator requires several mm, its large footprint is costly.

3.5.1 Conventional III−V modulators

EA modulators and MZMs have been used in relatively long transmission distance and high-capacity networks to overcome the limitations of DMLs. In addition, these devices are suitable for uncooled operation because the modulation speed is independent of the operation temperature. This is important because the removal of the thermoelectric cooler reduces the overall power consumption.

Fig. 3.16 shows the throughput and 3-dB bandwidth progression of EA modulators. Because the detuning of MQW is set small for IM, the EA modulator enables us to achieve large ER with several-hundred-μm active length. A 3-dB bandwidth of more than 40 GHz is achieved by using the conventional lumped electrode [91]. To further increase the bandwidth, an EA modulator using a traveling wave type of electrode has achieved a 3-dB bandwidth of more than 70 GHz [92]. As Table 3.5 shows, an EA modulator monolithically integrated with DFB lasers (externally modulated lasers (EMLs)) has been reported as a high-speed and small-chip-size optical light source [91,93−97] that can reduce the device cost.

Fig. 3.17 shows the throughput and 3-dB bandwidth change of MZMs. The MZMs generate optical signals by changing the refractive index using the QCSE and the Pockels effect. The length of the phase modulator is several mm because the detuning of the MQW is set to be large enough to avoid large absorption. Therefore the traveling-wave electrode is key to achieving a high modulation speed. For this purpose, there are two designs: a capacitance-loaded traveling-wave electrode (CL-TWE) and an n-i-p-n structure. Due to low resistance, MZMs with n-i-p-n CL-TWE have achieved a 3-dB bandwidth of more than 67 GHz, which is the measurement limit [102,103]. The MZMs can generate phase modulation format optical signals that include quadrature phase-shift keying and 16QAM [103−105].

Table 3.5 Comparison summary of vertical-cavity surface-emitting lasers (VCSELs), distributed feedback lasers, electro-absorption (EA) modulators, and Mach−Zehnder modulators (MZMs).

	VCSELs	**Directly modulated laser s**	**EA modulators**	**MZMs**
Schematic structure	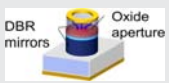	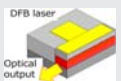	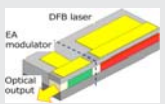	
Power consumption	Very low	Very low	Low	Little high
Optical output power	Very high	Very high	High	Little low
Chip size	Very small (0.1−0.3 mm × 0.1−0.3 mm)	Very small (0.2−0.4 mm × 0.15−0.2 mm)	Small (0.2−0.4 mm × 0.5−1.0 mm)	Large (0.5−2.5 mm × 3.0−7.0 mm)
3-dB bandwidth	∼30 GHz	<40 GHz	∼100 GHz (traveling wave type)	>67 GHz
ER	<6 dB	<6 dB	∼8 dB	∼10 dB
Chirp parameter α	∼3.0	∼3.0	0−10	0

ER, extinction ratio.

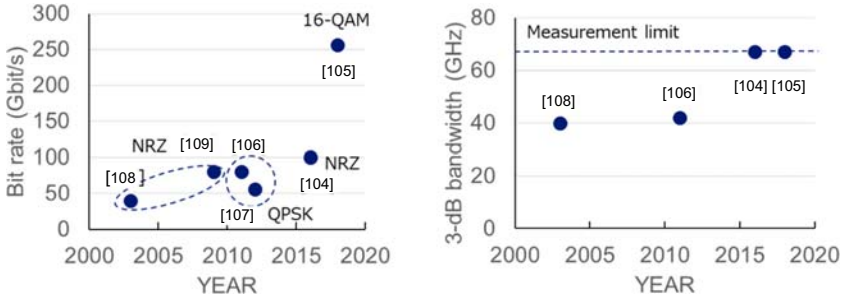

Figure 3.17 Bitrate and 3-dB bandwidth change of MZMs. [102–107]. *MZMs*, Mach–Zehnder modulator.

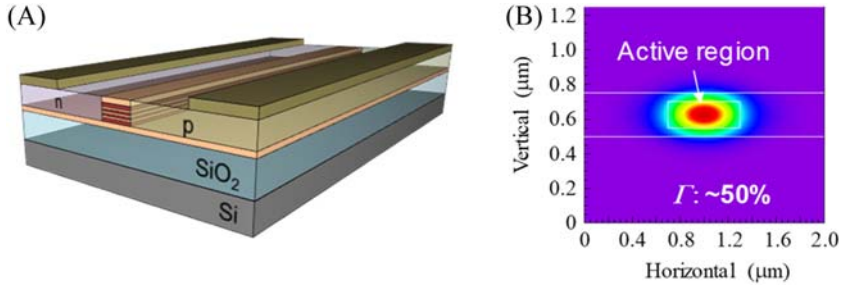

Figure 3.18 (A) Schematic diagram of membrane photonic device. (B) Calculated optical mode profile.

However, the disadvantages are the large chip size and the transmitter loss, which includes MZMs, a polarization convertor, and waveguides.

3.5.2 Membrane III–V modulators

The increase in optical confinement factors has been a focus in studies on achieving high modulation efficiency that results in low power consumption. Fig. 3.18A shows a schematic diagram of membrane III–V devices on SiO_2/Si substrate, in which a ~250-nm thick III–V layer with a lateral *pn* junction is sandwiched between low refractive index materials [108–113].

An optical confinement factor three times that of the conventional III–V modulator could be obtained from this structure, as shown in Fig. 3.18B. Utilizing a buried heterostructure can efficiently confine carriers and photons. Furthermore, integration with the Si waveguide is easy because the membrane III–V device has a size and effective refractive index similar to the Si waveguide. Fig. 3.19A shows a microscope image of a membrane MZM integrated with a membrane DFB laser [114]. The authors connected the laser and the phase modulator with a silicon 2×2 multimode interferometer. The active region length was 500 μm for both the laser

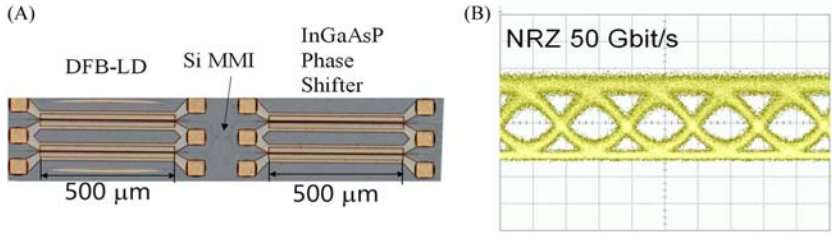

Figure 3.19 (A) Photograph of monolithically integrated MZM and DFB laser. (B) Eye diagram with 50 Gbit/s NRZ signal modulation. *DFB*, Distributed feedback; *MZM*, Mach−Zehnder modulator.

and the phase modulator. An n-doped bulk InGaAsP was used for the phase modulator, and the measured V_π ranged from 7.6 to 8.2 V, which corresponded to the $V_\pi L$ of around 0.4 V · cm. Due to the large optical confinement factor and the small capacitance of the membrane structure, 50-Gbit/s nonreturn-to-zero (NRZ) modulation performance using lumped electrode design was realized, as shown in Fig. 3.19B.

3.6 Heterogeneous integration on silicon

This section discusses recent work on combining a range of materials with the silicon photonics platform to exploit the benefits of heterogeneous integration. Each of the material integration strategies is discussed in turn.

3.6.1 III−V on silicon

As discussed earlier, the key metrics for a modulator are small size, low loss, low drive voltage, and high bandwidth. The first three metrics all improve with a larger modulation efficiency, and the last metric also depends on modulation efficiency because smaller devices have less microwave loss, resulting in larger bandwidths. Si modulators use the plasma effect to change the index and that works well and is widely used. III−V materials have a plasma effect, but also band filling, Kerr and Pockels effects, as shown in Fig. 3.20, their modulation efficiency tends to be about 10 times larger [115,116]. The plasma effect is orientation independent, but the Pockels effect is orientation dependent, so modulators must be properly oriented for the best performance. MZI modulators using III−V on Si have been demonstrated with bandwidths of 25 GHz and $V_\pi L$ of 2.4 V_{mm} [115]. Smaller, higher speed devices can be fabricated using EA [92]. Low power can be achieved with short devices operating at high impedance, but larger bandwidths typically require traveling wave designs with 50-Ω impedance as shown in Fig. 3.21, and bandwidths in excess of 67 GHz with modulation efficiency of 4.3 dB/V have been demonstrated [92].

The lowest power Si modulators utilize the MOSCAP design described earlier [36], but bonded III−V on Si MOS modulators have even lower drive voltages due

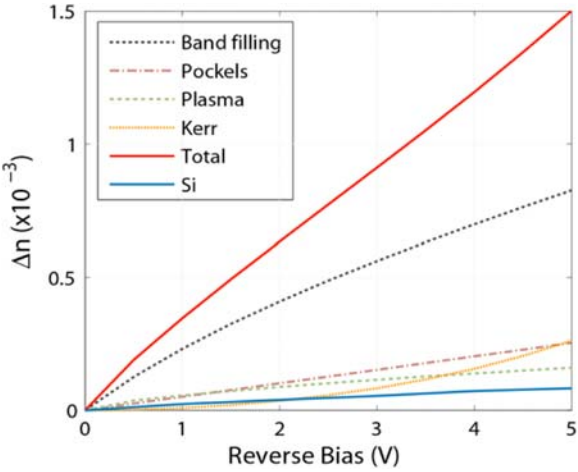

Figure 3.20 Index change per volt of Si (blue) and III−V (red) materials showing the relative contributions of bandfilling, plasma, Kerr, and Pockels effects for InGaAsP. *Source*: After H.-W. Chen, "High-Speed Hybrid Silicon Mach-Zehnder Modulator and Tunable Microwave Filter," University of California, Santa Barbara, 2011.

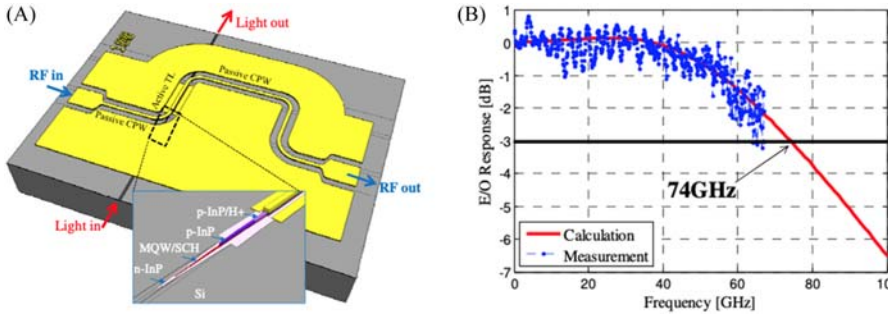

Figure 3.21 (A) Schematic diagram of an HSEAM. (B) Frequency response of the HSEAM modulator. *HSEAM*, Heterogeneous silicon electro-absorption modulator.

to the larger modulation efficiency of III−V materials. A schematic diagram of the ring version of this modulator is shown in Fig. 3.22 along with a transmission electron microscopy (TEM) cross-sectional image of the MOS capacitor. When this capacitor is charged or discharged by applying an electrical field between the III−V material and Si, the carrier concentration around the gate oxide can change by several orders of magnitude to exceed $10^{20}/cm^2$. Because this capacitor (dis) charge process occurs under a fast electric field effect, 40 Gb/s NRZ modulation is possible. However, compared with the pure Si design where deposited poly-Si is typically used as the top layer above the gate oxide, a wafer-bonded III−V layer offers three advantages with respect to its material characteristics: first, crystalline

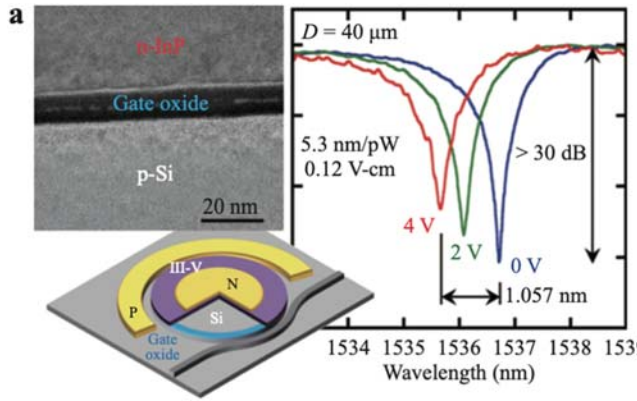

Figure 3.22 Schematic device structure and a cross-sectional TEM image of a III−V/SiO_2/Si MOS modulator. Also shown is the modulator loss for several voltages [119].
Source: After T. Hiraki et al., Heterogeneously integrated III-V/Si MOS capacitor Mach-Zehnder modulator, Nat. Photonics, 11 (8) (2017) 482−485.

III−V material has lower optical loss than poly-Si, even though it is highly doped. Second, the much smaller effective mass and larger mobility of electrons in a III−V material result in a larger real part change in the modal index, which translates into larger phase change efficiency and lower free-carrier absorption. Finally, the band-filling effect in a III−V material contributes to the phase change in addition to the plasma dispersion effect (Fig. 3.20). The $V_\pi L$ of a high-speed modulator is just 0.12 V · cm with a record tuning energy efficiency of 5.3 nm/pW. This is a huge improvement over conventional thermal and carrier injection tuning with a few or sub-nm/mW efficiency. An interesting configuration is to integrate this ring modulator with a gain region, resulting in a three-terminal configuration [117] with potentially all of the advantages of loss-modulated lasers rather than current-modulated lasers [118].

3.6.2 Organic electro-optic materials on Si

Integration of organic materials on silicon waveguides allows a combination of the tremendous processing advantages of silicon photonics with the wealth of optical properties that can be obtained by molecular engineering of organic materials [120]. Silicon-organic hybrid (SOH) devices have been demonstrated to provide second-order [121,122] and third-order [123] optical nonlinearities as well as optical gain [124,125]. For efficient devices, it is imperative to maximize the interaction of the optical waveguide mode with the organic cladding materials, for example, by exploiting discontinuous field enhancement in slot-waveguide structures [126,127].

The concept of a SOH MZM is sketched in Fig. 3.23A and B. Each arm of the modulator consists of a silicon photonic slot waveguide, comprising a pair of silicon rails. The rails are separated by a narrow slot with typical widths between 80 and

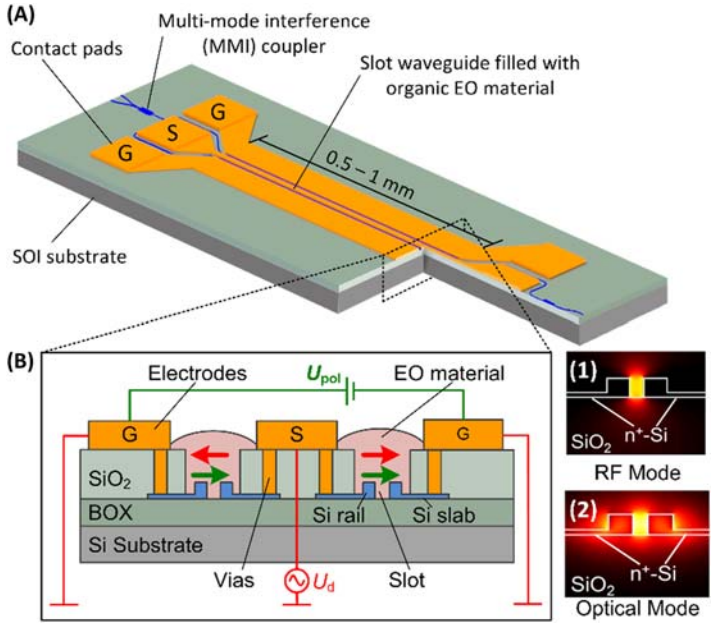

Figure 3.23 Device concept of a SOH EO MZM. (A) Layout of the SOH MZM: silicon photonic slot waveguides are embedded into an organic EO cladding. Light is split by an MMI coupler into the arms of the MZM. (B) Cross section of the SOH MZM: each slot waveguide consists of a pair of silicon photonic rails (width 200−240 nm, height ∼220 nm), separated by a slot (80−160 nm). The EO cladding homogeneously fills the slot. The signal is applied via a GSG coplanar transmission line, connected to the rails of the slot waveguide by conductive n-doped silicon slabs (height ∼70 nm) and metal vias. For poling of the EO materials, a DC voltage U_{pol} (*green circuit*) is applied across the floating ground electrodes at an elevated temperature. The molecules align accordingly and maintain their orientation (*green arrows*) after cooling. A modulating voltage U_d induces an electric field in the slot (*red arrows*) that is antiparallel in the left to the poling direction, and parallel to the poling direction in the right arm of the MZM, leading to push−pull operation. **(1)** Field distribution of the modulating electrical RF mode. **(2)** Field distribution of the optical quasi-TE mode. The strong overlap of the RF mode and the optical mode leads to highly efficient modulation [128]. *EO*, Electro-optic; *GSG*, ground-signal-ground; *MMI*, multimode interference; *MZM*, Mach−Zehnder modulator; *SOH*, silicon-organic hybrid.
Source: Adapted from O. Marshall, M. Hsu, Z. Wang, B. Kunert, C. Koos, D. Van Thourhout, Heterogeneous integration on silicon photonics, Proc. IEEE, 106 (12) (2018) 2258−2269.

160 nm, and the electrical RF drive signal is applied via a ground-signal-ground transmission line and copropagates with the optical signal along the device. The transmission line is connected to the silicon rails of the slot waveguide by metal vias and thin *n*-doped silicon slabs. The electrical RF voltage drops mainly across the narrow slot, to which the light is strongly confined, thus leading to a strong overlap of the optical mode and the modulating RF field, (see Insets (1) and (2) of

Fig. 3.23B). The $\chi^{(2)}$-nonlinearity of the organic EO material is activated in a one-time poling process by heating the device and by applying an external DC voltage U_{pol} to the floating ground electrodes, thereby aligning the EO molecules as indicated by the green arrows in Fig. 3.23B. The modulating fields (*red arrows*) associated with the RF drive voltage U_d are oriented parallel to the chromophore alignment in one arm and antiparallel in the other arm of the MZM, thereby leading to efficient chirp-free push−pull modulation.

The performance of SOH EO modulators has been demonstrated in a series of experiments, exhibiting remarkable efficiencies with voltage-length products down to $V_\pi L = 0.3$ Vmm and energy consumptions of only a few fJ/bit [129−132]. The viability of SOH devices has been shown both for data transmission using IM and direct detection (IM/DD) [133−135] and for coherent communications [131,132,136,137]. Regarding IM/DD transmission, generation of 100 Gbit/s OOK data signals [134] and 200 Gbit/s PAM-4 was shown [135]. For coherent communications, generation of advanced modulation formats such as 16QAM was demonstrated at ultralow energy consumption [131,132] and with symbol rates of up to 100 GBd transmitted on a single polarization [137]. The extraordinarily low drive voltage of SOH modulators allows operation of the devices directly from standard output ports of field-programable gate arrays without additional drive amplifiers [132]. Besides optical communications, SOH EO modulators have been used for frequency comb generation [138] or as ultracompact phase shifters for optical metrology [139]. SOH devices can be monolithically cointegrated with the full device portfolio of silicon photonics [132], thus offering a route toward highly functional silicon photonic circuits with advanced EO modulators.

The response of the EO cladding materials is ultrafast such that the bandwidth of SOH devices is typically limited by resistive-capacitive (RC) time constants that are associated with the slot capacitance and the resistivity of the silicon slabs [140]. The modulation speed can be greatly increased by plasmonic devices, which reach modulation bandwidths of hundreds of gigahertz [141,142], but suffer from high optical losses, quantified by π-voltage-loss products that are of the order of 10 V dB [122]. This trade-off can be overcome by using optimized doping profiles or [140] capacitive coupling between the RF transmission line and the rails of the optical slot waveguide [143].

Besides boosting the performance of SOH devices, research efforts have also concentrated on improved organic EO materials. In most material systems, less than 15% of the EO activity inherent in the chromophores is translated to macroscopic EO effects by poling. Improved theoretical methods [120] have led to new classes of organic EO materials and allowed systematic improvements of the poling efficiency. This resulted in macroscopic EO coefficients r_{33} in excess of 500 pm/V in thin films [144−148] and of up to 390 pm/V in SOH devices [130]. Theoretical calculations also suggest that new chromophores with significantly improved molecular first-order hyperpolarizability are possible [148], which would enable SOH modulators with π-voltages of less than 100 mV.

In view of industrial applications of SOH devices, systematic investigation and optimization of thermal and photochemical stability of EO cladding materials is of

utmost importance. First investigations have demonstrated long-term thermal stability in accordance with Telcordia standards for high-temperature storage [149], while photochemical stability is the subject of ongoing research. A specific advantage of organic EO materials is that a variety of parameters such as EO activity, optical loss, dielectric permittivity, lattice hardness, material compatibility, and material processability can be simultaneously optimized by systematic chemical modification of the material. SOH devices will hence immediately benefit from future advances in design and synthesis of functional organic optical materials.

3.6.3 Graphene on Si

Graphene is a gapless 2D material with broadband electronic and optoelectronic properties [150,151] and with conduction and valence bands in contact in the neutrality (Dirac) point. The Fermi level, which is related to the conductivity and carrier concentration, can be controlled by applying an external electric field. The gapless nature and the electrical control of carrier concentration lead to broadband EA and electro-refraction. When graphene is integrated into waveguides, efficient active photonic devices such as modulators, detectors, switches, and phase shifters can be realized. Graphene can be conveniently integrated on top of passive photonic platforms such as silicon (Si), silicon nitride (SiN), or other dielectric waveguides [152].

Graphene-based modulation exploits charge accumulation in a capacitor structure [153,154]. The first demonstration of an OOK modulation at 10 Gb/s was achieved with a single-layer graphene (SLG) placed on a thin oxide layer on top of a doped Si waveguide [155]. This device was operated at 1550 nm biased at 0.45 eV doping level and with a swing voltage that permitted the modulation of absorption. More recently, the same design was improved to 16 GHz bandwidth [156]. Larger bandwidth has been reported on SiN [157] and Si passive waveguides [158], up to 39 GHz EO bandwidth, 50 Gb/s data rate [159], and 2.2 dB/V modulation efficiency [160]. In all latter cases, the capacitor was made of two graphene layers spaced by a thin dielectric spacer, the so-called double-layer graphene (DLG). The DLG configuration features several advantages: it does not require carrier implantations or Ge epitaxy because modulation is achieved exclusively with graphene, the modulation efficiency is larger than in the SLG case thanks to the two graphene layers rather than one, and the IL is minimized because the waveguide is undoped. A key aspect of graphene modulators is the quality of graphene in terms of carrier mobility or equivalently in terms of carrier scattering time (Fig. 3.24).

High mobility, of the order of 10,000 $cm^2/V/s$ at 0.1 eV doping level (or 2200 $cm^2/V/s$ at 0.4 eV, that is a typical operating value), determines the minimum background loss of the material, that is, it influences fundamental FoM of modulators such as IL and ER [152]. For high doping levels, say above 0.55 eV at 1550 nm, the graphene response is mainly electro-refractive where phase modulation is dominant, and the modulation efficiency compared to typical Si photonics modulators is larger. A value of $V_\pi L = 0.28\ V \cdot cm$ measured in an SLG case has

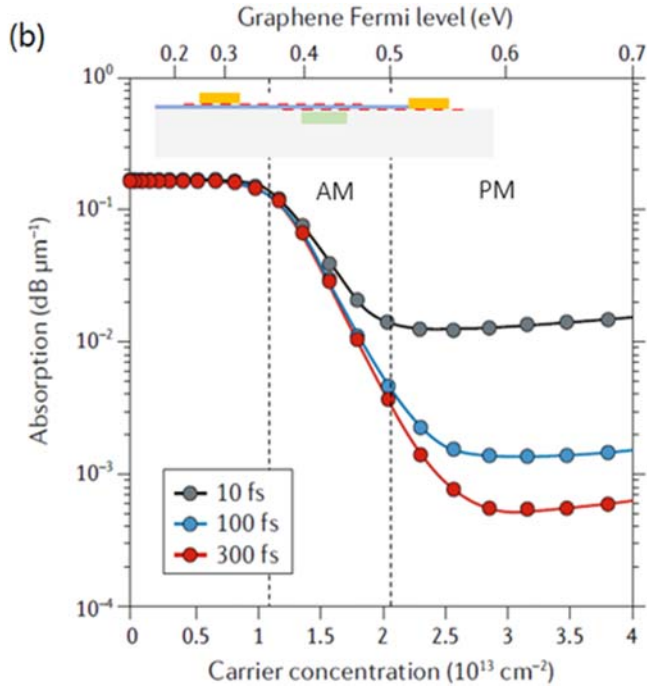

Figure 3.24 Optical absorption at 1550 nm of the DLG on SOI passive waveguide as a function of doping level for three values of scattering time constant (10fs, 100fs, 300fs) showing the effect of the carrier mobility on IL and ER. AM and PM indicate the biasing to operate the DLG in electro-absorption or electro-refraction mode [161]. *AM*, Amplitude modulation; *DLG*, Double-layer graphene; *ER*, extinction ratio; *IL*, insertion loss; *PM*, Phase Modulation; *SOI*, silicon on insulator.
Source: From V. Sorianello, G. Contestabile, M. Romagnoli, Graphene on silicon modulators, J. Light. Technol., 38 (10) (2020) 2782−2789.

been reported [162] whereas a value of $V_\pi L < 0.16\ \mathrm{V \cdot cm}$ for DLG has been predicted [152].

Another key parameter of the graphene modulator is the contact resistance. For high-speed modulators a convenient value is below 250-$\Omega \cdot$mm, which is a value achievable with both top contact and edge contact technologies [163,164].

Fabrication of graphene modulators requires wafer-scale chemical vapour deposition (CVD) graphene growth and transfer from the catalyst to the photonic wafer. This is mostly a postprocess done in the back end of line. Recently, wafer-scale fabrication has been developed in Si fabs [165] and the process can be compatible with CMOS foundries. This ensures the compatibility of graphene photonics in application-specific roadmaps such as those for optical interconnections and optical communications. Graphene modulators are promising in terms of bandwidth and modulation efficiency. State-of-the-art performances above 100 Gb/s are feasible.

3.6.4 Lithium niobate on Si

LN EO modulators, traditionally realized in bulk LN crystal, have been the workhorse of telecommunication for many decades, owing to their high bandwidth, linear response, and excellent material stability [166]. These legacy devices that rely on low-index contrast waveguides formed by ion indiffusion, or proton-exchange, are usually a few centimeters long and require driving voltage of several volts because of limited EO modulation efficiencies (Fig. 3.25A). The large size, discrete nature, high energy consumption, and high cost make the legacy system unsuitable for modern data communication applications.

Recent development in integrated thin-film LN (TFLN) photonics, enabled by the development of ultralow loss optical waveguides realized using conventional planar fabrication techniques [12], is promising a new class of devices with ultralow driving voltages (<1 V), wide bandwidth (>100 GHz) and low IL that can be tightly integrated in the same way as silicon photonics devices [167,168]. Achieving a low driving voltage could enable direct CMOS drive that reduces cost, complexity, and energy consumption. Linear and flat response could further reduce digital signal processing workload. The low IL is also key to bringing down total energy consumption in data communication systems. The simple design, proven material property, and excellent thermal stability offer an attractive solution to data communication applications. Recent key demonstrations include 700 Gbps intensity modulated-direct detection (IMDD) using single modulator [169], 1.55 Tbps coherent link at 200 Gbaud [170], and analog modulation up to 500 GHz [171].

The technical breakthrough that enabled these record demonstrations is the ability to fabricate high-confinement and low-loss waveguides in TFLN [172–174]. Achieving low-loss waveguides in TFLN (<0.3 dB/cm) allows the electrodes to be positioned closer to each other than their legacy designs resulting in large $V_\pi \cdot L_\pi$ improvement from about 10 V · cm at 1550 nm [175] in legacy modulators to about 2 V · cm at 1550 nm (1.6 V · cm at 1310 nm) in TFLN [167]. Such a large increase in modulation efficiency results in dramatically improved voltage and bandwidth performance and more dense integration. Two popular modulator integration approaches being explored are (1) directly patterned and etched TFLN-on-insulator

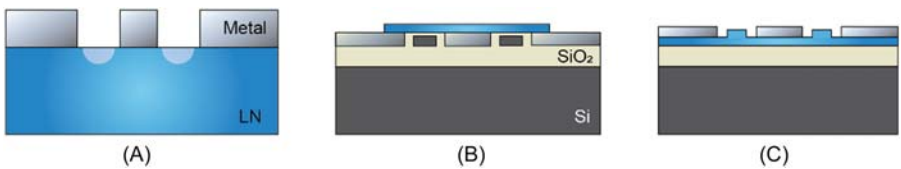

Figure 3.25 Cross section views of lithium niobate electro-optic modulators. (A) Traditional bulk lithium niobate with ion-indiffusion or proton-exchange waveguides, shown in light blue. The electrodes are thick and are separated by large gaps typically more than 10 μm. (B) Thin-film lithium niobate modulator with etched waveguides. The electrode gap can be significantly reduced to 5 μm or less. (C) Thin-film lithium niobate modulator based on back-end processing on a silicon-on-insulator platform. This approach has better compatibility to existing CMOS processes.

to create rib or channel waveguides [17,176] (Fig. 3.25B) and (2) backend film bonding integration allowing the optical mode to be evanescently coupled into a TFLN slab from silicon waveguides (e.g., Refs. [158] [Fig. 3.25C]). While the direct patterning approach has the potential to deliver the best possible modulator performance, the hybrid approach leverages the large library of developed silicon photonics components that can be integrated with TFLN.

The key technical challenge that remains for TFLN is a reliable and cost-effective interface between TFLN modulators, with micron-scale optical waveguides, and single-mode fibers. While the on-chip loss for TFLN modulators has been reduced to sub-dB levels, the fiber interface dominates the optical loss budget. Using concepts similar to those used in silicon photonics mode converters, progress has been made on mode transformation between discrete waveguide layers [177] and integrated edge couplers with lensed fibers, featuring coupling loss <1 dB/facet [178]. The open challenge is to find a high-yield mode converter design and manufacturing process that can be integrated with the silicon platform cost-effectively. This problem is further amplified by a manufacturing challenge: TFLN is not a CMOS-compatible material, and therefore TFLN-integrated photonics will likely require its own dedicated process lines. Early work has demonstrated wafer-scale manufacturability in TFLN-integrated systems, using standard CMOS tools [179] and compatibility with back-end processing in silicon photonics with the film bonding approach [158].

The large boost in performance combined with the ability to integrate modulators with other EO functionalities, such as frequency comb sources [180], on the same TFLN chip, could lead to system-level advantages, thus providing enough incentive to overcome the manufacturing-related capital expenditure. Fundamentally, the TFLN approach can be cost-effective because LN crystals, silicon wafers, and the bonding processes are all low cost and well established. We could envisage TFLN becoming an attractive solution for data communication applications due to its advantageous performances especially at ultrahigh speeds that next-generation data communication systems demand.

3.6.5 $BaTiO_3$ on Si

With Pockels coefficients above $r > 1000$ pm/V [181], BTO is one of the materials with the strongest Pockels effect known today [182]. Because the intrinsic EO material properties are 30 times larger than in LN and because it is a solid-state crystal with better thermal and optical stability than organics, BTO became an attractive material for integrated optical modulators. Challenges in BTO such as strong off-axis tensor coefficients, multidomain structure, and integration as thin films into a photonic platform have, however, only recently been solved. In this section, we summarize the history of BTO modulators, the progress on components on silicon, and the integration of BTO with silicon photonics.

The first demonstration of waveguide-based modulators using BTO was made on magnesium oxide (MgO) substrates around the year 2000 [183,184]. The substrates enabled epitaxial growth of films via metal-organic chemical vapor deposition and

pulsed laser deposition and ensured optical confinement in the BTO layer ($n \sim 2.3$) on top of MgO ($n \sim 1.7$). Epitaxial growth is required, first to obtain high-quality films with strong optical properties [185] and, second, to align the waveguides at specific angles relative to the BTO crystal, to exploit the strongest tensor elements [186]. The early work on MgO demonstrated the feasibility of realizing BTO-based integrated modulators, but the actual device performance was moderate ($V_\pi \sim 20-50$ V) due to large optical modes and limited material quality ($r \sim 50$ pm/V in Ref. [183]). In addition, the scalability to cost-efficient photonic platforms was unsolved due to the choice of substrates.

The development of high-quality, epitaxial BTO layers with strong Pockels coefficients on silicon ($r \sim 150$ pm/V) via molecular-beam epitaxy was reported in 2013 [187] (Fig. 3.26), and it enabled the integration of such high-Pockels material into silicon photonic platforms and reactivated research and development (R&D) activities on BTO modulators. Indeed, the first demonstrations of BTO-based passive waveguides [188] and modulators on silicon substrates followed quickly [189,190]. Waveguides were formed in an amorphous silicon layer on top of epitaxial BTO layers on silicon-on-insulator substrates, forming slot-waveguide structures (Fig. 3.27A). The modulator performance with $V_\pi L_\pi \sim 1.5$ V · cm exceeded many alternative modulator technologies, but bandwidth (<5 GHz) and losses (~ 50 dB/cm) represented limitations in that initial device generation. By adjusting process conditions and waveguide cross sections, the propagation losses were reduced to <3 dB/cm [191,192] in subsequent work. Indeed, moving away from high-capacitance slot-waveguide structures and optimizing mode overlap resulted in ring resonators with higher bandwidths of >30 GHz and better efficiency of $V_\pi L \sim 0.45$ V · cm [18] (Fig. 3.27B). Optical design optimizations were additionally supported by enhancing the properties of the films: bulk-like Pockels coefficients of $r > 900$ pm/V were reached in thin BTO films on silicon [18]. The field is working on further improvements of the device performance, in particular on reducing required voltages to <1 V, increasing the bandwidth to >100 GHz, and enabling broadband operation.

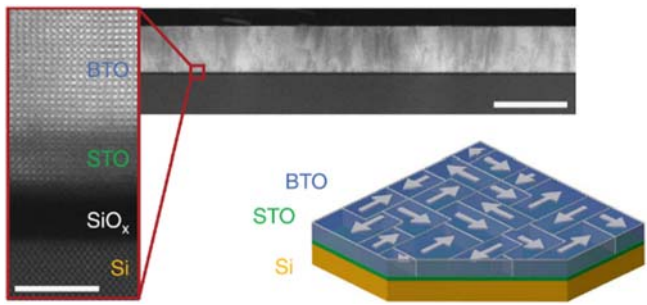

Figure 3.26 Fabrication of epitaxial BTO layers with strong Pockels effect directly on silicon wafers [187]. *BTO*, Barium titanate.
Source: From S. Abel et al., "A strong electro-optically active lead-free ferroelectric integrated on silicon," Nat. Commun., 4 (1) (2013) 1671.

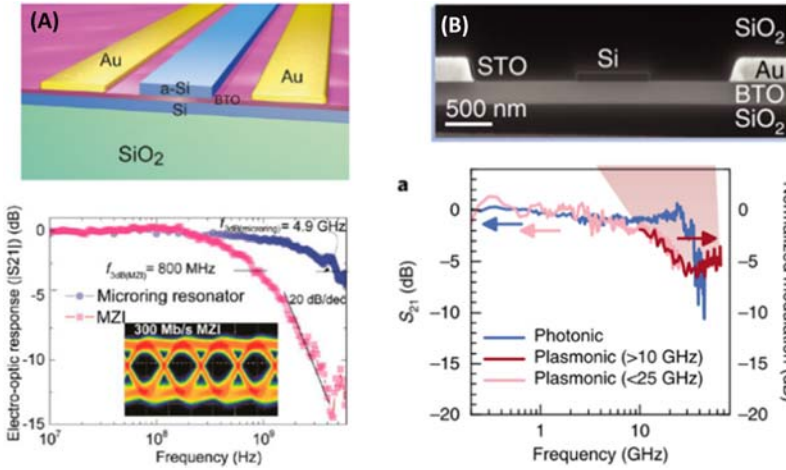

Figure 3.27 BTO-based optical modulators fabricated on silicon substrates with slot-waveguide (left [189]) and strip-loaded structures (right [18]). Please note that also the 3dB bandwidth of plasmonic modulators made from BTO is shown (see Ref. [18] for details).
BTO, Barium titanate.
Source: From (left) C. Xiong et al., Active silicon integrated nanophotonics: ferroelectric $BaTiO_3$ devices, Nano Lett., 14 (3) (2014) 1419−1425; (right) S. Abel et al., "Large Pockels effect in micro- and nanostructured barium titanate integrated on silicon," Nat. Mater., 18 (1) (2019) 42−47.

These improvements require a coherent optimization of device designs and material stacks, taking into account optical, electro-optical, and RF properties.

Progress in the performance of individual devices fostered the development of integration concepts of the BTO technology into a scalable photonic platform. A particular target is to combine BTO-based Pockels modulators with standard silicon photonics since other devices such as detectors, passives, and couplers are readily available. Various integration routes such as deposition of films during the front-end-of-line (FEOL) process are in principle possible. One particular integration scheme has already been demonstrated: the integration of BTO modulators onto silicon photonic wafers during the back-end-of-line process [193] (Fig. 3.28), with two important results: first, the performance of FEOL silicon photonic structures is maintained in that integration process, as shown in high-speed germanium detectors as the typically most sensitive photonic component. Second, BTO-based modulators are indeed functional after the integration, with excellent $V_\pi L \sim 0.2\ \mathrm{V} \cdot \mathrm{cm}$. Although BTO-integration schemes are not yet available in standard silicon photonic design and manufacturing routes, the first commercial development of BTO-based silicon photonics chips at scalable volumes has already been started [194].

Besides the target of reaching low-loss and high-speed Pockels modulators for communication, the strong EO effect has strong potential for other upcoming applications—in particular when available in a scalable, low-cost photonic platform. Examples of such new applications are nonvolatile optical memory elements [195],

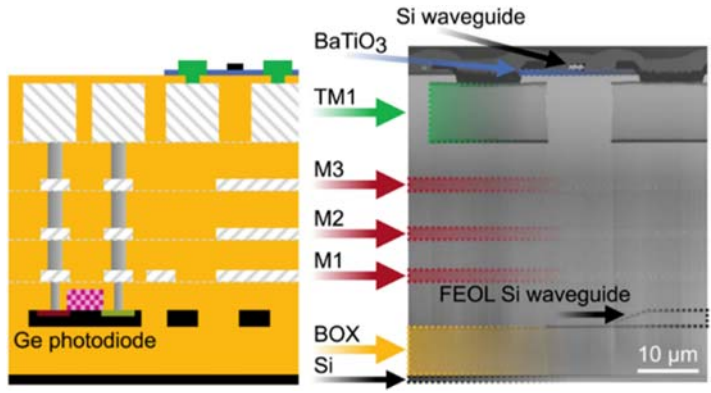

Figure 3.28 Example of the integration of BTO-based photonic modulators onto a silicon photonic platform. *BTO*, Barium titanate.
Source: From F. Eltes et al., A $BaTiO_3$-based electro-optic Pockels modulator monolithically integrated on an advanced silicon photonics platform, J. Light. Technol., 37 (5) (2019) 1456–1462.

ultracompact Pockels-based plasmonic modulators [18,196] (Fig. 3.27B), low-power tuning elements with three orders better power efficiency than traditional Joule heaters in silicon photonics [197], and efficient low-loss high-speed modulators operating at cryogenic temperatures for quantum transducers [192]. The rapid development of the photonic BTO technology with strong improvements in device performance within only a few years sets a promising outlook for having available a silicon photonic platform empowered with an ultrastrong Pockels effect.

3.7 Summary and perspectives

Silicon or silicon-compatible optical modulator is arguably the most researched of all silicon photonic devices. It has a storied 35-year history starting with Richard Soref and Brian Bennett's proposal in 1987 to use the FCPD effect for phase modulation in crystalline silicon [5]. Numerous silicon modulator designs were demonstrated in the years that followed, but they were limited in speed until the MOSCAP-based free-carrier dispersion device in 2004 showed >1 GHz bandwidth [36]. This seminal demonstration played a big part in igniting the explosion of silicon photonic R&D and investment across academia and industry. The result is a vibrant field of study with designs fueling optical transceivers in volume production now as well as designs promising orders of magnitude improvement in size, bandwidth density, and power consumption for future applications.

Today, the most proven high-speed silicon photonic modulators use the FCPD effect of reverse bias p-n diodes in the two branches of MZI designs. They support 25.78 Gbps NRZ, 53.125 Gbps PAM4, and 106.25 Gbps PAM4 modulation of datacom's optical pluggable multi source agreement (MSA) products. Similar designs

are being developed for the next speed node of 212.5 Gbps PAM4. Besides speed scaling, bandwidth density and power are key performance drivers as the industry moves to pack more data into the same pluggable module form factors and, more aggressively, into much smaller copackage optics footprints, as discussed in Chapters 6 and 7. Ring resonator-based devices have the needed reduction in size, and designs using reverse bias p-n diodes or MOS capacitors for high-speed phase modulation have demonstrated good performance. So they are among the most promising modulators for next-generation datacom applications. Ge FKE and QCSE modulators also have the needed compact form factor, but they must demonstrate improved modulation efficiency in the O-band, which is the standard for datacom, and mature their material growth onto the silicon photonic platform before they can be viable for product.

The exponential growth of data, which is doubling every 2−3 years, demands similar rapid scaling of bandwidth density and power. To meet future performance targets, high-speed optical modulation requires disruptive innovation. As a result, there is a large body of research focused on integrating onto the silicon photonic platform materials with much better EO coefficients than silicon and germanium. III−V bonding onto Si to create a MOS capacitor device for modulation is a promising approach because of its demonstrated performance and its bonding, and integration process is similar to what has already been proven for hybrid laser integration, see Chapter 2. Integrating other materials like EO polymers, graphene, LN, $BaTiO_3$, etc. onto the silicon platform is also interesting because they have the potential to enable much more efficient and therefore more compact and lower power modulators than what is possible with native silicon designs. Nevertheless, their adoption for product will be critically dependent on our ability to overcome the complexity and cost of hybrid integration. While it is not clear which hybrid approach will ultimately surpass and displace proven silicon devices, the high-interest level and steady progress in this area underscore researchers' commitment to aggressively scale performance and are reasons for us to be hopeful that additional technology breakthroughs are in our near future.

References

[1] L. Vivien, L. Pavesi, vol. Chapter 9 Handbook of Silicon Photonics, Taylor & Francis, 2016.

[2] G. Cocorullo, F.G. Della Corte, I. Rendina, Temperature dependence of the thermo-optic coefficient in crystalline silicon between room temperature and 550 K at the wavelength of 1523 nm, Appl. Phys. Lett. 74 (22) (1999) 3338−3340. Available from: https://doi.org/10.1063/1.123337.

[3] D. Perez-Galacho, et al., Integrated polarization beam splitter for 100/400 GE polarization multiplexed coherent optical communications, J. Light. Technol. 32 (3) (2014) 361−368. Available from: https://doi.org/10.1109/JLT.2013.2293782.

[4] G.T. Reed (Ed.), Silicon Photonics: The State of the Art, John Wiley & Sons, Ltd, Chichester, UK, 2008.

[5] R. Soref, B. Bennett, Electrooptical effects in silicon, IEEE J. Quantum Electron. 23 (1) (1987) 123−129. Available from: https://doi.org/10.1109/JQE.1987.1073206.
[6] A. Frova, P. Handler, F.A. Germano, D.E. Aspnes, Electro-absorption effects at the band edges of silicon and germanium, Phys. Rev. 145 (2) (1966) 575−583. Available from: https://doi.org/10.1103/PhysRev.145.575.
[7] G.T. Reed, et al., Recent breakthroughs in carrier depletion based silicon optical modulators, Nanophotonics 3 (4−5) (2014) 229−245. Available from: https://doi.org/10.1515/nanoph-2013-0016.
[8] F.Y. Gardes, D.J. Thomson, N.G. Emerson, G.T. Reed, 40 Gb/s silicon photonics modulator for TE and TM polarisations, Opt. Express 19 (12) (2011) 11804. Available from: https://doi.org/10.1364/OE.19.011804.
[9] D. Marris-Morini, et al., Low loss 40 Gbit/s silicon modulator based on interleaved junctions and fabricated on 300 mm SOI wafers, Opt. Express 21 (19) (2013) 22471. Available from: https://doi.org/10.1364/OE.21.022471.
[10] Q. Xu, S. Manipatruni, B. Schmidt, J. Shakya, M. Lipson, 125 Gbit/s carrier-injection-based silicon micro-ring silicon modulators, Opt. Express 15 (2) (2007) 430. Available from: https://doi.org/10.1364/OE.15.000430.
[11] I. Charlet, Y. Desieres, D. Marris-Morini, F. Boeuf, Capacitive modulator design optimization using Si and strained-SiGe for datacom applications, IEEE J. Sel. Top. Quantum Electron. 27 (3) (2021) 1−8. Available from: https://doi.org/10.1109/JSTQE.2020.3028447.
[12] D. Galacho, High Speed Optical Modulation, Advanced Modulation Formats and Mode Division Multiplexing in Silicon Photonics, Université Paris-Saclay, 2016.
[13] L. Vivien et al., "Building Blocks of Silicon Photonics," Semiconductors and Semimetals, Elsevier, 101 (2019) 1−41.
[14] L. Mastronardi, et al., High-speed Si/GeSi hetero-structure electro absorption modulator, Opt. Express 26 (6) (2018) 6663. Available from: https://doi.org/10.1364/OE.26.006663.
[15] P. Chaisakul, et al., 23 GHz Ge/SiGe multiple quantum well electro-absorption modulator, Opt. Express 20 (3) (2012) 3219. Available from: https://doi.org/10.1364/OE.20.003219.
[16] M.-S. Rouifed, et al., Advances toward Ge/SiGe quantum-well waveguide modulators at 1.3 μm, IEEE J. Sel. Top. Quantum Electron. 20 (4) (2014) 33−39. Available from: https://doi.org/10.1109/JSTQE.2013.2294456.
[17] C. Wang, et al., Integrated lithium niobate electro-optic modulators operating at CMOS-compatible voltages, Nature 562 (7725) (2018) 101−104. Available from: https://doi.org/10.1038/s41586-018-0551-y.
[18] S. Abel, et al., Large Pockels effect in micro- and nanostructured barium titanate integrated on silicon, Nat. Mater. 18 (1) (2019) 42−47. Available from: https://doi.org/10.1038/s41563-018-0208-0.
[19] M. Berciano, et al., Fast linear electro-optic effect in a centrosymmetric semiconductor, Commun. Phys. 1 (1) (2018) 64. Available from: https://doi.org/10.1038/s42005-018-0064-x.
[20] D. Munk, et al., Surface acoustic wave photonic devices in silicon on insulator, Nat. Commun. 10 (1) (2019) 4214. Available from: https://doi.org/10.1038/s41467-019-12157-x.
[21] B.S. Dennis, M.I. Haftel, D.A. Czaplewski, D. Lopez, G. Blumberg, V.A. Aksyuk, Compact nanomechanical plasmonic phase modulators, Nat. Photonics 9 (4) (2015) 267−273. Available from: https://doi.org/10.1038/nphoton.2015.40.
[22] R.A. Soref, B.R. Bennett, "Kramers-Kronig Analysis of Electro-Optical Switching in Silicon," Mar. 1987, pp. 32. Available from: https://doi.org/10.1117/12.937193.

[23] J.P. Lorenzo, R.A. Soref, 1.3 μm electro-optic silicon switch, Appl. Phys. Lett. 51 (1) (1987) 6–8. Available from: https://doi.org/10.1063/1.98887.
[24] L. Friedman, R.A. Soref, J.P. Lorenzo, Silicon double-injection electro-optic modulator with junction gate control, J. Appl. Phys. 63 (6) (1988) 1831–1839. Available from: https://doi.org/10.1063/1.339878.
[25] C.K. Tang, G.T. Reed, A.J. Walton, A.G. Rickman, Simulation of a low loss optical modulator for fabrication in Simox material, MRS Proc. 298 (1993) 247. Available from: https://doi.org/10.1557/PROC-298-247.
[26] C.K. Tang, G.T. Reed, A.J. Walton, A.G. Rickman, Low-loss, single-model optical phase modulator in SIMOX material, J. Light. Technol. 12 (8) (1994) 1394–1400. Available from: https://doi.org/10.1109/50.317527.
[27] G.T. Reed, C.E. Jason Png, Silicon optical modulators, Mater. Today 8 (1) (2005) 40–50. Available from: https://doi.org/10.1016/S1369-7021(04)00678-9.
[28] A. Cutolo, M. Iodice, P. Spirito, L. Zeni, Silicon electro-optic modulator based on a three terminal device integrated in a low-loss single-mode SOI waveguide, J. Light. Technol. 15 (3) (1997) 505–518. Available from: https://doi.org/10.1109/50.557567.
[29] C.K. Tang, G.T. Reed, Highly efficient optical phase modulator in SOI waveguides, Electron. Lett. 31 (6) (1995) 451–452. Available from: https://doi.org/10.1049/el:19950328.
[30] S.M. Jackson, et al., Optical beamsteering using integrated optical modulators, J. Light. Technol. 15 (12) (1997) 2259–2263. Available from: https://doi.org/10.1109/50.643551.
[31] S.M. Jackson, et al., A novel optical phase modulator design suitable for phased arrays, J. Light. Technol. 16 (11) (1998) 2016–2019. Available from: https://doi.org/10.1109/50.730364.
[32] P.D. Hewitt, G.T. Reed, Improving the response of optical phase modulators in SOI by computer simulation, J. Light. Technol. 18 (3) (2000) 443–450. Available from: https://doi.org/10.1109/50.827519.
[33] P.D. Hewitt, G.T. Reed, Improved modulation performance of a silicon p-i-n device by trench isolation, J. Light. Technol. 19 (3) (2001) 387–390. Available from: https://doi.org/10.1109/50.918892.
[34] I. Day et al., "Tapered silicon waveguides for low insertion loss highly-efficient high-speed electronic variable optical attenuators," in: *Optical Fiber Communications Conference*, 2003, pp. 249–251, Available from: https://doi.org/10.1109/OFC.2003.1247634.
[35] P. Dainesi, et al., CMOS compatible fully integrated Mach-Zehnder interferometer in SOI technology, IEEE Photonics Technol. Lett. 12 (6) (2000) 660–662. Available from: https://doi.org/10.1109/68.849076.
[36] A. Liu, et al., A high-speed silicon optical modulator based on a metal–oxide–semiconductor capacitor, Nature 427 (6975) (2004) 615–618. Available from: https://doi.org/10.1038/nature02310.
[37] L. Liao, et al., High speed silicon Mach-Zehnder modulator, Opt. Express 13 (8) (2005) 3129. Available from: https://doi.org/10.1364/OPEX.13.003129.
[38] C.E. Png, S.P. Chan, S.T. Lim, G.T. Reed, Optical phase modulators for MHz and GHz modulation in silicon-on-insulator (SOI), J. Light. Technol. 22 (6) (2004) 1573–1582. Available from: https://doi.org/10.1109/JLT.2004.827655.
[39] Q. Xu, B. Schmidt, S. Pradhan, M. Lipson, Micrometre-scale silicon electro-optic modulator, Nature 435 (7040) (2005) 325–327. Available from: https://doi.org/10.1038/nature03569.
[40] F.Y. Gardes, G.T. Reed, N.G. Emerson, C.E. Png, A sub-micron depletion-type photonic modulator in silicon on insulator, Opt. Express 13 (22) (2005) 8845. Available from: https://doi.org/10.1364/OPEX.13.008845.

[41] A. Liu, et al., High-speed optical modulation based on carrier depletion in a silicon waveguide, Opt. Express 15 (2) (2007) 660. Available from: https://doi.org/10.1364/OE.15.000660.
[42] L. Liao, et al., 40 Gbit/s silicon optical modulator for high-speed applications, Electron. Lett. 43 (22) (2007) 1196. Available from: https://doi.org/10.1049/el:20072253.
[43] D.J. Thomson, et al., High contrast 40Gbit/s optical modulation in silicon, Opt. Express 19 (12) (2011) 11507. Available from: https://doi.org/10.1364/OE.19.011507.
[44] D.J. Thomson, et al., 50-Gb/s silicon optical modulator, IEEE Photonics Technol. Lett. 24 (4) (2012) 234−236. Available from: https://doi.org/10.1109/LPT.2011.2177081.
[45] K. Li, et al., Electronic−photonic convergence for silicon photonics transmitters beyond 100 Gbps on−off keying, Optica 7 (11) (2020) 1514. Available from: https://doi.org/10.1364/OPTICA.411122.
[46] M. Chagnon, et al., Design, analysis, and transmission system performance of a 41 GHz silicon photonic modulator, Opt. Express 23 (11) (2015) 14263. Available from: https://doi.org/10.1364/oe.23.014263.
[47] S. Meister, et al., Matching p-i-n-junctions and optical modes enables fast and ultra-small silicon modulators, Opt. Express 21 (13) (2013) 16210. Available from: https://doi.org/10.1364/OE.21.016210.
[48] S. Akiyama, et al., Compact PIN-diode-based silicon modulator using side-wall-grating waveguide, IEEE J. Sel. Top. Quantum Electron. 19 (6) (2013) 74−84. Available from: https://doi.org/10.1109/JSTQE.2013.2278438.
[49] A. Shakoor, K. Nozaki, E. Kuramochi, K. Nishiguchi, A. Shinya, M. Notomi, Compact 1D-silicon photonic crystal electro-optic modulator operating with ultra-low switching voltage and energy, Opt. Express 22 (23) (2014) 28623. Available from: https://doi.org/10.1364/OE.22.028623.
[50] Y. Sobu, G. Huang, S. Tanaka, Y. Tanaka, Y. Akiyama, T. Hoshida, High-speed optical digital-to-analog converter operation of compact two-segment all-silicon Mach−Zehnder modulator, J. Light. Technol. 39 (4) (2021) 1148−1154. Available from: https://doi.org/10.1109/JLT.2020.3035684.
[51] J. Van Campenhout, M. Pantouvaki, P. Verheyen, S. Selvaraja, G. Lepage, H. Yu, W. Lee, J. Wouters, D. Goossens, M. Moelants, W. Bogaerts, P. Absil, Low-voltage, low-loss, multi-Gb/s silicon micro-ring modulator based on a MOS capacitor, in: *OFC/NFOEC*, 2012, pp. 1−3.
[52] M. Webster et al., "An efficient MOS-capacitor based silicon modulator and CMOS drivers for optical transmitters," in: *11th International Conference on Group IV Photonics (GFP)*, Aug. 2014, pp. 1−2, Available from: https://doi.org/10.1109/JLT.2020.3035684.
[53] M. Sodagar, et al., Compact, 15 Gb/s electro-optic modulator through carrier accumulation in a hybrid $Si/SiO_2/Si$ microdisk, Opt. Express 23 (22) (2015) 28306. Available from: https://doi.org/10.1364/OE.23.028306.
[54] K. Debnath, et al., All-silicon carrier accumulation modulator based on a lateral metal-oxide-semiconductor capacitor, Photonics Res. 6 (5) (2018) 373. Available from: https://doi.org/10.1364/PRJ.6.000373.
[55] W. Zhang, et al., High bandwidth capacitance efficient silicon MOS modulator, J. Light. Technol. 39 (1) (2021) 201−207. Available from: https://doi.org/10.1109/JLT.2020.3026945.
[56] M. Sakib et al., "A high-speed micro-ring modulator for next generation energy-efficient optical networks beyond 100 Gbaud," in: *Conference on Lasers and Electro-Optics*, 2021, pp. SF1C.3. Available from: https://doi.org/10.1364/CLEO_SI.2021.SF1C.3.

[57] W. Franz, Einfluß eines elektrischen Feldes auf eine optische Absorptionskante, Z. für Naturforsch. A 13 (6) (1958) 484–489. Available from: https://doi.org/10.1515/zna-1958-0609.
[58] L.V. Keldysh, Behavior of non-metallic crystals in strong electric fields, Sov. J. Exp. Theor. Phys. 6 (1958) 7631958, Accessed: Oct. 09, 2021. [Online]. Available. Available from: https://ui.adsabs.harvard.edu/abs/1958JETP . . ..6.0.763K/abstract.
[59] S. Jongthammanurak, et al., Large electro-optic effect in tensile strained Ge-on-Si films, Appl. Phys. Lett. 89 (16) (2006) 161115. Available from: https://doi.org/10.1063/1.2363948.
[60] C. Porret et al., "O-band GeSi quantum-confined stark effect electro-absorption modulator integrated in a 220nm silicon photonics platform," in: *2020 IEEE Symposium on VLSI Technology*, Jun. 2020, pp. 1–2. Available from: https://doi.org/10.1109/VLSITechnology18217.2020.9265082.
[61] P. Chaisakul et al., "High speed electro-absorption modulator based on quantum-confined stark effect from Ge/SiGe multiple quantum wells," in: *The 9th International Conference on Group IV Photonics (GFP)*, Aug. 2012, pp. 60–62. Available from: https://doi.org/10.1109/GROUP4.2012.6324086.
[62] E.H. Edwards, et al., Low-voltage broad-band electroabsorption from thin Ge/SiGe quantum wells epitaxially grown on silicon, Opt. Express 21 (1) (2013) 867. Available from: https://doi.org/10.1364/OE.21.000867.
[63] J.E. Roth, et al., C-band side-entry Ge quantum-well electroabsorption modulator on SOI operating at 1 V swing, Electron. Lett. 44 (1) (2008) 49. Available from: https://doi.org/10.1049/el:20082979.
[64] S. Ren, et al., Ge/SiGe quantum well waveguide modulator monolithically integrated with SOI waveguides, IEEE Photonics Technol. Lett. 24 (6) (2012) 461–463. Available from: https://doi.org/10.1109/LPT.2011.2181496.
[65] K. Zang, et al., Germanium quantum well QCSE waveguide modulator with tapered coupling in distributed modulator–detector system, J. Light. Technol. 35 (21) (2017) 4629–4633. Available from: https://doi.org/10.1109/JLT.2017.2753582.
[66] L. Mastronardi, GeSi Franz-Keldysh modulator for silicon photonic integrated circuits (Doctoral thesis), University of Southampton, 2019. <https://scholar.google.co.uk/scholar?hl = en&as_sdt = 0%2C5&q = GeSi + Franz-Keldysh + modulator + for + silicon + photonic + integrated + circuits&btnG = > (accessed 09.10.21).
[67] Y.-H. Kuo, et al., Quantum-confined stark effect in Ge/SiGe quantum wells on Si for optical modulators, IEEE J. Sel. Top. Quantum Electron. 12 (6) (2006) 1503–1513. Available from: https://doi.org/10.1109/JSTQE.2006.883146.
[68] A.E.-J. Lim, et al., Novel evanescent-coupled germanium electro-absorption modulator featuring monolithic integration with germanium p-i-n photodetector, Opt. Express 19 (6) (2011) 5040. Available from: https://doi.org/10.1364/OE.19.005040.
[69] S.A. Srinivasan, et al., 56 Gb/s germanium waveguide electro-absorption modulator, J. Light. Technol. 34 (2) (2016) 419–424. Available from: https://doi.org/10.1109/JLT.2015.2478601.
[70] J. Liu, et al., Waveguide-integrated, ultralow-energy GeSi electro-absorption modulators, Nat. Photonics 2 (7) (2008) 433–437. Available from: https://doi.org/10.1038/nphoton.2008.99.
[71] D. Feng, et al., High-speed GeSi electroabsorption modulator on the SOI waveguide platform, IEEE J. Sel. Top. Quantum Electron. 19 (6) (2013) 64–73. Available from: https://doi.org/10.1109/JSTQE.2013.2278881.
[72] A.V. Krishnamoorthy, et al., A low-power, high-speed, 9-channel germanium-silicon electro-absorption modulator array integrated with digital CMOS driver and wavelength

multiplexer, Opt. Express 22 (10) (2014) 12289−12295. Available from: https://doi.org/10.1364/OE.22.012289.

[73] O. El-Aassar, G.M. Rebeiz, A DC-to-108-GHz CMOS SOI distributed power amplifier and modulator driver leveraging multi-drive complementary stacked cells, IEEE J. Solid-State Circuits 54 (12) (2019) 3437−3451. Available from: https://doi.org/10.1109/JSSC.2019.2941013.

[74] Y. Zhao, L. Vera, J.R. Long, A 10 Gb/s, 6 V p-p, digitally controlled, differential distributed amplifier MZM driver, IEEE J. Solid-State Circuits 49 (9) (2014) 2030−2043. Available from: https://doi.org/10.1109/JSSC.2014.2327036.

[75] J. Kim, J.F. Buckwalter, A 40-Gb/s optical transceiver front-end in 45 nm SOI CMOS, IEEE J. Solid-State Circuits 47 (3) (2012) 615−626. Available from: https://doi.org/10.1109/JSSC.2011.2178723.

[76] D.-U. Li, C.-M. Tsai, 10-Gb/s modulator drivers with local feedback networks, IEEE J. Solid-State Circuits 41 (5) (2006) 1025−1030. Available from: https://doi.org/10.1109/JSSC.2006.872878.

[77] H. Li, et al., A 112 Gb/s PAM4 silicon photonics transmitter with microring modulator and CMOS driver, J. Light. Technol. 38 (1) (2020) 131−138. Available from: https://doi.org/10.1109/JLT.2019.2938731.

[78] C. Xiong, D.M. Gill, J.E. Proesel, J.S. Orcutt, W. Haensch, W.M.J. Green, Monolithic 56 Gb/s silicon photonic pulse-amplitude modulation transmitter, Optica 3 (10) (2016) 1060. Available from: https://doi.org/10.1364/OPTICA.3.001060.

[79] Q. Liao et al., "A 50Gb/s high-efficiency Si-photonic transmitter with lump-segmented MZM and integrated PAM4 CDR," in: *2021 IEEE Custom Integrated Circuits Conference (CICC)*, Apr. 2021, pp. 1−2. Available from: https://doi.org/10.1109/CICC51472.2021.9431546.

[80] H. Li, et al., A 3-D-integrated silicon photonic microring-based 112-Gb/s PAM-4 transmitter with nonlinear equalization and thermal control, IEEE J. Solid-State Circuits 56 (1) (2021) 19−29. Available from: https://doi.org/10.1109/JSSC.2020.3022851.

[81] M. Kim, D.-H. Kwon, D.-W. Rho, W.-Y. Choi, A low-power 28-Gb/s PAM-4MZM driver with level pre-distortion, IEEE Trans. Circuits Syst. II Express Briefs 68 (3) (2021) 908−912. Available from: https://doi.org/10.1109/TCSII.2020.3020128.

[82] C. Li, et al., Silicon photonic transceiver circuits with microring resonator bias-based wavelength stabilization in 65 nm CMOS, IEEE J. Solid-State Circuits 49 (6) (2014) 1419−1436. Available from: https://doi.org/10.1109/JSSC.2014.2321574.

[83] S. Shopov, S.P. Voinigescu, A 3×60Gb/s transmitter/repeater front-end with 4.3VPP single-ended output swing in a 28nm UTBB FD-SOI technology, IEEE J. Solid-State Circuits 51 (7) (2016) 1651−1662. Available from: https://doi.org/10.1109/JSSC.2016.2545703.

[84] N. Qi, et al., Co-design and demonstration of a 25-Gb/s silicon-photonic Mach−Zehnder modulator with a CMOS-based high-swing driver, IEEE J. Sel. Top. Quantum Electron. 22 (6) (2016) 131−140. Available from: https://doi.org/10.1109/JSTQE.2016.2602102.

[85] T. Jyo, M. Nagatani, J. Ozaki, M. Ishikawa, H. Nosaka, 12.3 A 48GHz BW 225mW/ch linear driver IC with stacked current-reuse architecture in 65nm CMOS for beyond-400Gb/s coherent optical transmitters, in: *2020 IEEE International Solid- State Circuits Conference - (ISSCC)*, 2020, pp. 212−214. Available from: https://doi.org/10.1109/ISSCC19947.2020.9063027.

[86] K. Li et al., "Co-design of electronics and photonics components for silicon photonics transmitters," in: *2018 European Conference on Optical Communication (ECOC)*, Sep. 2018, pp. 1−3. Available from: https://doi.org/10.1109/ECOC.2018.8535212.

[87] D. Petousi, et al., Monolithically integrated high-extinction-ratio MZM with a segmented driver in photonic BiCMOS, IEEE Photonics Technol. Lett. 28 (24) (2016) 2866−2869. Available from: https://doi.org/10.1109/LPT.2016.2624700.
[88] T.N. Huynh, et al., Flexible transmitter employing silicon-segmented Mach−Zehnder modulator with 32-nm CMOS distributed driver, J. Light. Technol. 34 (22) (2016) 5129−5136. Available from: https://doi.org/10.1109/JLT.2016.2606558.
[89] H. Sepehrian, A. Yekani, L.A. Rusch, W. Shi, CMOS-photonics codesign of an integrated DAC-Less PAM-4 silicon photonic transmitter, IEEE Trans. Circuits Syst. I Regul. Pap. 63 (12) (2016) 2158−2168. Available from: https://doi.org/10.1109/TCSI.2016.2613514.
[90] P. Rito, et al., A monolithically integrated segmented linear driver and modulator in EPIC 0.25- μ m SiGe:C BiCMOS platform, IEEE Trans. Microw. Theory Tech. 64 (12) (2016) 4561−4572. Available from: https://doi.org/10.1109/TMTT.2016.2618392.
[91] Ch. Jany, Ch. Kazmierski, J. Decobert, F. Alexandre, F. Blache, O. Drisse, D. Carpentier, N. Lagay, F. Martin, E. Derouin, T. Johansen, Ch. Jiang, Semi-insulating buried heterostructure 1.55 micrometer InGaAlAs electroabsorption modulated laser with 60GHz bandwidth, in: *33rd European Conference and Exhibition of Optical Communication − post-deadline papers*, 16−20 September 2007.
[92] Y. Tang, J.D. Peters, J.E. Bowers, Over 67 GHz bandwidth hybrid silicon electroabsorption modulator with asymmetric segmented electrode for 13 μm transmission, Opt. Express 20 (10) (2012) 11529. Available from: https://doi.org/10.1364/OE.20.011529.
[93] M.R. Gokhale, P.V. Studenkov, J. Ueng-McHale, J. Thomson, J. Yao, Uncooled, 10Gb/s 1310 nm electroabsorption modulated laser, Opt. Fiber Commun. Conf. Opt. Soc. Am. (2021). p. PD42, 2003, Accessed: Oct. 10. Available from: https://www.osapublishing.org/abstract.cfm?uri = ofc-2003-PD42.
[94] W. Kobayashi, et al., Wide temperature range (-25 degrees C-100 degrees C) operation of a 10-Gb/s 1.55-mu m electroabsorption modulator integrated DFB laser for 80-km SMF transmission, IEEE Photonics Technol. Lett. 21 (15) (2009) 1054−1056. Available from: https://doi.org/10.1109/LPT.2009.2022182.
[95] H. Mardoyan et al., "204-GBaud on-off keying transmitter for inter-data center communications," in: *Optical Fiber Communication Conference Postdeadline Papers*, 2018, p. Th4A.4. Available from: https://doi.org/10.1364/OFC.2018.Th4A.4.
[96] K. Zhong, et al., 140-Gb/s 20-km transmission of PAM-4 signal at 1.3 μm for short reach communications, IEEE Photonics Technol. Lett. 27 (16) (2015) 1757−1760. Available from: https://doi.org/10.1109/LPT.2015.2439571.
[97] U. Troppenz et al., "1.3 μm electroabsorption modulated lasers for PAM4/PAM8 single channel 100 Gb/s," in: *26th International Conference on Indium Phosphide and Related Materials (IPRM)*, 2014, pp. 1−2. Available from: https://doi.org/10.1109/ICIPRM.2014.6880562.
[98] W. Kobayashi, et al., Novel approach for chirp and output power compensation applied to a 40-Gbit/s EADFB laser integrated with a short SOA, Opt. Express 23 (7) (2015) 9533. Available from: https://doi.org/10.1364/OE.23.009533.
[99] S. Kanazawa, et al., Flip-chip interconnection lumped-electrode EADFB laser for 100-Gb/s/ λ transmitter, IEEE Photonics Technol. Lett. 27 (16) (2015) 1699−1701. Available from: https://doi.org/10.1109/LPT.2015.2438076.
[100] S. Kanazawa, et al., 214-Gb/s 4-PAM operation of flip-chip interconnection EADFB Laser module, J. Light. Technol. 35 (3) (2017) 418−422. Available from: https://doi.org/10.1109/JLT.2016.2632164.

[101] X. Wu et al., A 20Gb/s NRZ/PAM-4 1V transmitter in 40nm CMOS driving a Si-photonic modulator in 0.13 μm CMOS, in: *2013 IEEE International Solid-State Circuits Conference Digest of Technical Papers*, Feb. 2013, pp. 128−129. Available from: https://doi.org/10.1109/ISSCC.2013.6487667.

[102] Y. Ogiso, et al., Over 67 GHz bandwidth and 1.5 V V_{π} InP-based optical IQ modulator with n-i-p-n heterostructure, J. Light. Technol. 35 (8) (2017) 1450−1455. Available from: https://doi.org/10.1109/JLT.2016.2639542.

[103] Y. Ogiso et al., "Ultra-high bandwidth InP IQ modulators for next generation coherent transmitter," in: *2017 IEEE Compound Semiconductor Integrated Circuit Symposium (CSICS)*, Oct. 2017, pp. 1−4. Available from: https://doi.org/10.1109/CSICS.2017.8240423.

[104] K. Prosyk et al., "Tunable InP-based optical IQ modulator for 160 Gb/s," in: *37th European Conference and Exposition on Optical Communications*, 2011, pp. Th.13.A.5. Available from: https://doi.org/10.1364/ECOC.2011.Th.13.A.5.

[105] E. Yamada et al., "112-Gb/s InP DP-QPSK modulator integrated with a Silica-PLC polarization multiplexing circuit," in: *Proc. OFC, National Fiber Optic Engineers Conference*, 2012, pp. PDP5A.9. Available from: https://doi.org/10.1364/NFOEC.2012.PDP5A.9.

[106] K. Tsuzuki, et al., 40 Gbit/s n−i−n InP Mach−Zehnder modulator with a voltage of 2.2 V, Electron. Lett. 39 (20) (2003) 1464. Available from: https://doi.org/10.1049/el:20030939.

[107] F. Tsuzuki, T. Ishibashi, T. Ito, N. Kikuchi, F. Kano, 80 Gb/s InP Mach-Zehnder modulator module using liquid crystal polymer (LCP) transmission nine, in: *35th European Conference on Optical Communication*, 20−24 September 2009.

[108] S. Matsuo, T. Fujii, K. Hasebe, K. Takeda, T. Sato, T. Kakitsuka, Directly modulated buried heterostructure DFB laser on SiO_2/Si substrate fabricated by regrowth of InP using bonded active layer, Opt. Express 22 (10) (2014) 12139. Available from: https://doi.org/10.1364/OE.22.012139.

[109] S. Matsuo, T. Fujii, K. Hasebe, K. Takeda, T. Sato, T. Kakitsuka, Directly modulated DFB Laser on SiO_2/Si substrate for datacenter networks, J. Light. Technol. 33 (6) (2015) 1217−1222. Available from: https://doi.org/10.1109/JLT.2014.2386875.

[110] S. Matsuo, T. Kakitsuka, Low-operating-energy directly modulated lasers for short-distance optical interconnects, Adv. Opt. Photonics 10 (3) (2018) 567. Available from: https://doi.org/10.1364/AOP.10.000567.

[111] T. Fujii, et al., Heterogeneously integrated membrane lasers on Si substrate for low operating energy optical links, IEEE J. Sel. Top. Quantum Electron. 24 (1) (2018) 1−8. Available from: https://doi.org/10.1109/JSTQE.2017.2778510.

[112] H. Nishi, et al., Integration of eight-channel directly modulated membrane-laser array and SiN AWG multiplexer on Si, J. Light. Technol. 37 (2) (2019) 266−273. Available from: https://doi.org/10.1109/JLT.2018.2873742.

[113] T. Fujii, et al., Development of an epitaxial growth technique using III-V on a Si platform for heterogeneous integration of membrane photonic devices on Si, Appl. Sci. 11 (4) (2021) 1801. Available from: https://doi.org/10.3390/app11041801.

[114] T. Hiraki, et al., Integration of a high-efficiency Mach-Zehnder modulator with a DFB laser using membrane InP-based devices on a Si photonics platform, Opt. Express 29 (2) (2021) 2431. Available from: https://doi.org/10.1364/OE.411483.

[115] H.-W. Chen, J.D. Peters, J.E. Bowers, Forty Gb/s hybrid silicon Mach-Zehnder modulator with low chirp, Opt. Express 19 (2) (2011) 1455. Available from: https://doi.org/10.1364/OE.19.001455.

[116] H.-W. Chen, High-Speed Hybrid Silicon Mach-Zehnder Modulator and Tunable Microwave Filter, University of California, Santa Barbara, 2011.
[117] D. Liang, J.E. Bowers, Recent progress in heterogeneous III-V-on-silicon photonic integration, Light. Adv. Manuf. 2 (1) (2021) 1–25. Available from: https://doi.org/10.37188/lam.2021.005.
[118] D. Liang, X. Huang, G. Kurczveil, M. Fiorentino, R.G. Beausoleil, Integrated finely tunable microring laser on silicon, Nat. Photonics 10 (11) (2016) 719–722. Available from: https://doi.org/10.1038/nphoton.2016.163.
[119] T. Hiraki, et al., Heterogeneously integrated III–V/Si MOS capacitor Mach–Zehnder modulator, Nat. Photonics 11 (8) (2017) 482–485. Available from: https://doi.org/10.1038/nphoton.2017.120.
[120] L. Dalton, S. Benight, Theory-guided design of organic electro-optic materials and devices, Polym. (Basel) 3 (3) (2011) 1325–1351. Available from: https://doi.org/10.3390/polym3031325.
[121] T. Baehr-Jones, et al., Optical modulation and detection in slotted silicon waveguides, Opt. Express 13 (14) (2005) 5216. Available from: https://doi.org/10.1364/OPEX.13.005216.
[122] C. Koos, et al., Silicon-organic hybrid (SOH) and plasmonic-organic hybrid (POH) integration, J. Light. Technol. 34 (2) (2016) 256–268. Available from: https://doi.org/10.1109/JLT.2015.2499763.
[123] C. Koos, et al., All-optical high-speed signal processing with silicon–organic hybrid slot waveguides, Nat. Photonics 3 (4) (2009) 216–219. Available from: https://doi.org/10.1038/nphoton.2009.25.
[124] D. Korn, et al., Lasing in silicon–organic hybrid waveguides, Nat. Commun. 7 (1) (2016) 10864. Available from: https://doi.org/10.1038/ncomms10864.
[125] D. Kohler, et al., Biophotonic sensors with integrated Si3N4-organic hybrid (SiNOH) lasers for point-of-care diagnostics, Light. Sci. Appl. 10 (1) (2021) 64. Available from: https://doi.org/10.1038/s41377-021-00486-w.
[126] V.R. Almeida, Q. Xu, C.A. Barrios, M. Lipson, Guiding and confining light in void nanostructure, Opt. Lett. 29 (11) (2004) 1209. Available from: https://doi.org/10.1364/OL.29.001209.
[127] C. Koos, L. Jacome, C. Poulton, J. Leuthold, W. Freude, Nonlinear silicon-on-insulator waveguides for all-optical signal processing, Opt. Express 15 (10) (2007) 5976. Available from: https://doi.org/10.1364/OE.15.005976.
[128] O. Marshall, M. Hsu, Z. Wang, B. Kunert, C. Koos, D. Van Thourhout, Heterogeneous integration on silicon photonics, Proc. IEEE 106 (12) (2018) 2258–2269. Available from: https://doi.org/10.1109/JPROC.2018.2858542.
[129] S. Koeber, et al., Femtojoule electro-optic modulation using a silicon–organic hybrid device, Light. Sci. Appl. 4 (2) (2015) e255. Available from: https://doi.org/10.1038/lsa.2015.28. e255.
[130] C. Kieninger, et al., Ultra-high electro-optic activity demonstrated in a silicon-organic hybrid modulator, Optica 5 (6) (2018) 739. Available from: https://doi.org/10.1364/OPTICA.5.000739.
[131] M. Lauermann, et al., Low-power silicon-organic hybrid (SOH) modulators for advanced modulation formats, Opt. Express 22 (24) (2014) 29927. Available from: https://doi.org/10.1364/OE.22.029927.
[132] S. Wolf, et al., DAC-Less amplifier-less generation and transmission of QAM signals using sub-volt silicon-organic hybrid modulators, J. Light. Technol. 33 (7) (2015) 1425–1432. Available from: https://doi.org/10.1109/JLT.2015.2394511.

[133] H. Zwickel, et al., Silicon-organic hybrid (SOH) modulators for intensity-modulation / direct-detection links with line rates of up to 120 Gbit/s, Opt. Express 25 (20) (2017) 23784. Available from: https://doi.org/10.1364/OE.25.023784.
[134] S. Wolf, et al., Silicon-organic hybrid (SOH) Mach-Zehnder modulators for 100 Gbit/s on-off Keying, Sci. Rep. 8 (1) (2018) 2598. Available from: https://doi.org/10.1038/s41598-017-19061-8.
[135] C. Kieninger, et al., Silicon-organic hybrid (SOH) Mach-Zehnder modulators for 100 GBd PAM4 signaling with sub-1 dB phase-shifter loss, Opt. Express 28 (17) (2020) 24693. Available from: https://doi.org/10.1364/OE.390315.
[136] M. Lauermann, et al., 40 GBd 16QAM Signaling at 160 Gb/s in a silicon-organic hybrid modulator, J. Light. Technol. 33 (6) (2015) 1210−1216. Available from: https://doi.org/10.1109/JLT.2015.2394211.
[137] S. Wolf, et al., Coherent modulation up to 100 GBd 16QAM using silicon-organic hybrid (SOH) devices, Opt. Express 26 (1) (2018) 220. Available from: https://doi.org/10.1364/OE.26.000220.
[138] C. Weimann, et al., Silicon-organic hybrid (SOH) frequency comb sources for terabit/s data transmission, Opt. Express 22 (3) (2014) 3629. Available from: https://doi.org/10.1364/OE.22.003629.
[139] M. Lauermann, et al., Integrated optical frequency shifter in silicon-organic hybrid (SOH) technology, Opt. Express 24 (11) (2016) 11694. Available from: https://doi.org/10.1364/OE.24.011694.
[140] H. Zwickel, et al., Verified equivalent-circuit model for slot-waveguide modulators, Opt. Express 28 (9) (2020) 12951. Available from: https://doi.org/10.1364/OE.383120.
[141] S. Ummethala, et al., THz-to-optical conversion in wireless communications using an ultra-broadband plasmonic modulator, Nat. Photonics 13 (8) (2019) 519−524. Available from: https://doi.org/10.1038/s41566-019-0475-6.
[142] M. Burla, et al., 500 GHz plasmonic Mach-Zehnder modulator enabling sub-THz microwave photonics, APL Photonics 4 (5) (2019) 056106. Available from: https://doi.org/10.1063/1.5086868.
[143] S. Ummethala, et al., Hybrid electro-optic modulator combining silicon photonic slot waveguides with high-k radio-frequency slotlines, Optica 8 (4) (2021) 511. Available from: https://doi.org/10.1364/OPTICA.411161.
[144] L. Dalton, M. Lauermann, C. Koos, in: Fundamental Aspects of Materials and Applications, 2, World Scienific, 2016, pp. 369−396.
[145] P.A. G. J. Sullivan, L.R. Dalton, D.L. Andrews, The materials genome for organic electro-optics and silicon/plasmonic—organic hybrid technologyin N. Horiz. Nanosci. Eng. (2015) 233−284.
[146] L.E. Johnson, L.R. Dalton, B.H. Robinson, Optimizing calculations of electronic excitations and relative hyperpolarizabilities of electrooptic chromophores, Acc. Chem. Res. 47 (11) (2014) 3258−3265. Available from: https://doi.org/10.1021/ar5000727.
[147] L.R. Dalton, P. Gunter, M. Jazbinsek, O.-P. Kwon, P.A. Sullivan, Organic Electro-Optics and Photonics, Cambridge University Press, Cambridge, 2015.
[148] W. Jin, et al., Benzocyclobutene barrier layer for suppressing conductance in nonlinear optical devices during electric field poling, Appl. Phys. Lett. 104 (24) (2014) 243304. Available from: https://doi.org/10.1063/1.4884829.
[149] C. Kieninger, et al., Demonstration of long-term thermally stable silicon-organic hybrid modulators at 85°C, Opt. Express 26 (21) (2018) 27955. Available from: https://doi.org/10.1364/OE.26.027955.

[150] A.K. Bodenmann, A.H. MacDonald, Graphene: exploring carbon flatland, Phys. Today 60 (8) (2007) 35–41. Available from: https://doi.org/10.1063/1.2774096.
[151] F. Bonaccorso, Z. Sun, T. Hasan, A.C. Ferrari, Graphene photonics and optoelectronics, Nat. Photonics 4 (9) (2010) 611–622. Available from: https://doi.org/10.1038/nphoton.2010.186.
[152] M. Romagnoli, et al., Graphene-based integrated photonics for next-generation datacom and telecom, Nat. Rev. Mater. 3 (10) (2018) 392–414. Available from: https://doi.org/10.1038/s41578-018-0040-9.
[153] M. Liu, et al., A graphene-based broadband optical modulator, Nature 474 (7349) (2011) 64–67. Available from: https://doi.org/10.1038/nature10067.
[154] M. Liu, X. Yin, X. Zhang, Double-layer graphene optical modulator, Nano Lett. 12 (3) (2012) 1482–1485. Available from: https://doi.org/10.1021/nl204202k.
[155] Y. Hu, et al., Broadband 10 Gb/s operation of graphene electro-absorption modulator on silicon, Laser Photon. Rev. 10 (2) (2016) 307–316. Available from: https://doi.org/10.1002/lpor.201500250.
[156] C. Alessandri, et al., High speed graphene-silicon electro-absorption modulators for the O-band and C-band, Jpn. J. Appl. Phys. 59 (5) (2020) 052008. Available from: https://doi.org/10.35848/1347-4065/ab8920.
[157] C.T. Phare, Y.-H. Daniel Lee, J. Cardenas, M. Lipson, Graphene electro-optic modulator with 30 GHz bandwidth, Nat. Photonics 9 (8) (2015) 511–514. Available from: https://doi.org/10.1038/nphoton.2015.122.
[158] R. Safian, M. Teng, L. Zhuang, S. Chakravarty, Foundry-compatible thin film lithium niobate modulator with RF electrodes buried inside the silicon oxide layer of the SOI wafer, Opt. Express 28 (18) (2020) 25843. Available from: https://doi.org/10.1364/OE.396335.
[159] M.A. Giambra, et al., High-speed double layer graphene electro-absorption modulator on SOI waveguide, Opt. Express 27 (15) (2019) 20145. Available from: https://doi.org/10.1364/OE.27.020145.
[160] H. Agarwal, et al., 2D-3D integration of hexagonal boron nitride and a high-κ dielectric for ultrafast graphene-based electro-absorption modulators, Nat. Commun. 12 (1) (2021) 1070. Available from: https://doi.org/10.1038/s41467-021-20926-w.
[161] V. Sorianello, G. Contestabile, M. Romagnoli, Graphene on silicon modulators, J. Light. Technol. 38 (10) (2020) 2782–2789. Available from: https://doi.org/10.1109/JLT.2020.2974189.
[162] V. Sorianello, et al., Graphene–silicon phase modulators with gigahertz bandwidth, Nat. Photonics 12 (1) (2018) 40–44. Available from: https://doi.org/10.1038/s41566-017-0071-6.
[163] T. Cusati, et al., Electrical properties of graphene-metal contacts, Sci. Rep. 7 (1) (2017) 5109. Available from: https://doi.org/10.1038/s41598-017-05069-7.
[164] H.-Y. Park, et al., Extremely low contact resistance on graphene through n-type doping and edge contact design, Adv. Mater. 28 (5) (2016) 864–870. Available from: https://doi.org/10.1002/adma.201503715.
[165] M.A. Giambra, et al., Wafer-scale integration of graphene-based photonic devices, ACS Nano 15 (2) (2021) 3171–3187. Available from: https://doi.org/10.1021/acsnano.0c09758.
[166] E.L. Wooten, et al., A review of lithium niobate modulators for fiber-optic communications systems, IEEE J. Sel. Top. Quantum Electron. 6 (1) (2000) 69–82. Available from: https://doi.org/10.1109/2944.826874.
[167] M. Zhang, C. Wang, P. Kharel, D. Zhu, M. Lončar, Integrated lithium niobate electro-optic modulators: when performance meets scalability, Optica 8 (5) (2021) 652. Available from: https://doi.org/10.1364/OPTICA.415762.

[168] A. Rao, S. Fathpour, Compact lithium niobate electrooptic modulators, IEEE J. Sel. Top. Quantum Electron. 24 (4) (2018) 1−14. Available from: https://doi.org/10.1109/JSTQE.2017.2779869.

[169] X. Chen et al., "Single-wavelength and single-photodiode 700 Gb/s entropy-loaded PS-256-QAM and 200-GBaud PS-PAM-16 Transmission over 10-km SMF," in: *2020 European Conference on Optical Communications (ECOC)*, Dec. 2020, pp. 1−4. Available from: https://doi.org/10.1109/ECOC48923.2020.9333201.

[170] Chen, G. Raybon, D. Che, J. Cho, K. W. Kim, "Transmission of 200-GBaud PDM probabilistically shaped 64-QAM signals modulated via a 100-GHz thin-film LiNbO3 I/Q modulator, in: *2021 Optical Fiber Communications Conference*, 2021, pp. 1−3.

[171] A.J. Mercante, S. Shi, P. Yao, L. Xie, R.M. Weikle, D.W. Prather, Thin film lithium niobate electro-optic modulator with terahertz operating bandwidth, Opt. Express 26 (11) (2018) 14810. Available from: https://doi.org/10.1364/OE.26.014810.

[172] M. Zhang, C. Wang, R. Cheng, A. Shams-Ansari, M. Lončar, Monolithic ultra-high-Q lithium niobate microring resonator, Optica 4 (12) (2017) 1536. Available from: https://doi.org/10.1364/OPTICA.4.001536.

[173] I. Krasnokutska, J.-L.J. Tambasco, X. Li, A. Peruzzo, Ultra-low loss photonic circuits in lithium niobate on insulator, Opt. Express 26 (2) (2018) 897. Available from: https://doi.org/10.1364/OE.26.000897.

[174] R. Wu, et al., Lithium niobate micro-disk resonators of quality factors above 10 7, Opt. Lett. 43 (17) (2018) 4116. Available from: https://doi.org/10.1364/OL.43.004116.

[175] M. Doi, M. Sugiyama, K. Tanaka, M. Kawai, Advanced LiNbO/sub 3/optical modulators for broadband optical communications, IEEE J. Sel. Top. Quantum Electron. 12 (4) (2006) 745−750. Available from: https://doi.org/10.1109/JSTQE.2006.876192.

[176] M. Xu, et al., High-performance coherent optical modulators based on thin-film lithium niobate platform, Nat. Commun. 11 (1) (2020) 3911. Available from: https://doi.org/10.1038/s41467-020-17806-0.

[177] N. Boynton, et al., A heterogeneously integrated silicon photonic/lithium niobate travelling wave electro-optic modulator, Opt. Express 28 (2) (2020) 1868. Available from: https://doi.org/10.1364/OE.28.001868.

[178] P. Ying, et al., Low-loss edge-coupling thin-film lithium niobate modulator with an efficient phase shifter, Opt. Lett. 46 (6) (2021) 1478. Available from: https://doi.org/10.1364/OL.418996.

[179] K. Luke, P. Kharel, C. Reimer, L. He, M. Loncar, M. Zhang, Wafer-scale low-loss lithium niobate photonic integrated circuits, Opt. Express 28 (17) (2020) 24452. Available from: https://doi.org/10.1364/OE.401959.

[180] M. Xu, M. He, X. Cai, Generation of flat optical frequency comb using integrated cascaded lithium niobate modulators, in: *Conference on Lasers and Electro-Optics*, 2020, p. STh1O.5. Available from: https://doi.org/10.1364/CLEO_SI.2020.STh1O.5.

[181] A.R. Johnston, J.M. Weingart, Determination of the low-frequency linear electro-optic effect in tetragonal $BaTiO_3$†, J. Opt. Soc. Am. 55 (7) (1965) 828. Available from: https://doi.org/10.1364/JOSA.55.000828.

[182] S. Abel, J. Fompeyrine, Electro-optically active oxides on silicon for photonics, in Thin Films on Silicon (Material Energy), 8, World Scientific, Singapore, 2016, pp. 455−501.

[183] D.M. Gill, C.W. Conrad, G. Ford, B.W. Wessels, S.T. Ho, Thin-film channel waveguide electro-optic modulator in epitaxial $BaTiO_3$, Appl. Phys. Lett. 71 (13) (1997) 1783−1785. Available from: https://doi.org/10.1063/1.119397.

[184] A. Petraru, J. Schubert, M. Schmid, C. Buchal, Ferroelectric $BaTiO_3$ thin-film optical waveguide modulators, Appl. Phys. Lett. 81 (8) (2002) 1375−1377. Available from: https://doi.org/10.1063/1.1498151.
[185] K.J. Kormondy, et al., Microstructure and ferroelectricity of $BaTiO_3$ thin films on Si for integrated photonics, Nanotechnology 28 (7) (2017) 075706. Available from: https://doi.org/10.1088/1361-6528/aa53c2.
[186] P. Castera, et al., Electro-optical modulation based on Pockels effect in $BaTiO_3$ with a multi-domain structure, IEEE Photonics Technol. Lett. 28 (9) (2016) 990−993. Available from: https://doi.org/10.1109/LPT.2016.2522509.
[187] S. Abel, et al., A strong electro-optically active lead-free ferroelectric integrated on silicon, Nat. Commun. 4 (1) (2013) 1671. Available from: https://doi.org/10.1038/ncomms2695.
[188] S. Abel et al., "Electro-optical properties of barium titanate films epitaxially grown on silicon," 2012, pp. 82630Y. Available from: https://doi.org/10.1117/12.908772.
[189] C. Xiong, et al., Active silicon integrated nanophotonics: ferroelectric $BaTiO_3$ devices, Nano Lett. 14 (3) (2014) 1419−1425. Available from: https://doi.org/10.1021/nl404513p.
[190] S. Abel et al., "Barium Titanate Thin Films Embedded in Silicon Photonic Structures for Ultra-Efficient Electro-Optical Tuning," Sep. 2016. Available from: https://doi.org/10.7567/SSDM.2016.C-5-01.
[191] F. Eltes, et al., Low-loss $BaTiO_3$−Si waveguides for nonlinear integrated photonics, ACS Photonics 3 (9) (2016) 1698−1703. Available from: https://doi.org/10.1021/acsphotonics.6b00350.
[192] F. Eltes, et al., An integrated optical modulator operating at cryogenic temperatures, Nat. Mater. 19 (11) (2020) 1164−1168. Available from: https://doi.org/10.1038/s41563-020-0725-5.
[193] F. Eltes, et al., A $BaTiO_3$-based electro-optic Pockels modulator monolithically integrated on an advanced silicon photonics platform, J. Light. Technol. 37 (5) (2019) 1456−1462. Available from: https://doi.org/10.1109/JLT.2019.2893500.
[194] <http://www.lumiphase.com>.
[195] S. Abel, D.J. Stark, F. Eltes, J.E. Ortmann, D. Caimi, J. Fompeyrine, "Multi-level optical weights in integrated circuits," in: 2017 *IEEE International Conference on Rebooting Computing (ICRC)*, Nov. 2017, pp. 1−3. Available from: https://doi.org/10.1109/ICRC.2017.81236722.
[196] A. Messner, et al., Plasmonic ferroelectric modulators, J. Light. Technol. 37 (2) (2019) 281−290. Available from: https://doi.org/10.1109/JLT.2018.2881332.
[197] J.E. Ortmann, et al., Ultra-low-power tuning in hybrid barium titanate−silicon nitride electro-optic devices on silicon, ACS Photonics 6 (11) (2019) 2677−2684. Available from: https://doi.org/10.1021/acsphotonics.9b00558.

This page intentionally left blank

High-speed photodetectors

4

Kazumi Wada[1], Jifeng Liu[2], Yasuhiko Ishikawa[3], Donghwan Ahn[4], Dong Pan[5], Pengfei Cai[5] and Lionel Kimerling[6]
[1]Department of Materials Science and Engineering, Massachusetts Institute of Technology, Cambridge, MA, United States, [2]Materials Science and Engineering Program Lead, Thayer School of Engineering, Dartmouth College, Hanover, NH, United States, [3]Department of Electrical & Electronic Information Engineering, Toyohashi University of Technology, Tempaku Toyohashi, Japan, [4]School of Advanced Materials Engineering, College of Engineering, Kookmin University, Seoul, Korea, [5]SiFotonics Technologies Co. Ltd, Woburrn, MA, United States, [6]MIT Department of Materials Science and Engineering, Massachusetts Institute of Technology, Cambridge, MA, United States

4.1 Photodetectors in data center architecture

Monolithic electronic−photonic integration has the scaling potential to become the transistor of the 21st century. By providing a pathway for enormous bandwidth at all levels of the interconnection hierarchy, distributed systems with special purpose resources can reduce functional latency under moderate cost and energy requirements. The silicon platform provides the materials, tools, and design experience curve to define gaps and solutions for common, cross-market applications. This potential is perhaps most evident with the commercial advent of waveguide-integrated photodetectors (PDs). These devices are small, high speed, and, in the case of germanium absorbers, compatible with silicon materials' integration and foundry process design standards.

Optical cables today account for most of >100-m reach data center links. Furthermore, On-Board Optical traces are targeted for the remaining <3-m rack/board/package links. In-Package Optics now represents a critical opportunity for <20-mm chip/interposer links within data centers. Crossing each of these interconnect hierarchy thresholds presents huge manufacturing cost reduction requirements, to compete with the incumbent copper trace technology at a constant system-level cost.

The PD has been the driver for advancing Ge_xSi_{1-x} on Si integration at $x > 0.8$ [1,2]. The data center standard photon spectrum for optical interconnects is the O-band, centered at $\lambda = 1310$ nm. At this wavelength, Ge is essentially a direct gap semiconductor with an absorption length of ~ 1 μm, and optical fibers exhibit high transparency with loss ~ 0.4 dB/km. This combination of silicon, germanium, and silica glass has provided a direct path for detector waveguide integration and the

Integrated Photonics for Data Communication Applications. DOI: https://doi.org/10.1016/B978-0-323-91224-2.00013-8
© 2023 Elsevier Inc. All rights reserved.

system level benefits that it provides. Design rules are emerging for PD device structure, materials growth, and materials processing, as elaborated in this chapter. To strategically tune PD design, subject to the unique systems needs of data centers (e.g., requisite photonic link power budget, signal modulation scheme, latency), the modern photonic-integrated circuit (PIC) designer must develop proficiency in understanding the interplay of device structure, materials growth, and processing.

4.1.1 Performance parameters and waveguide–detector integration paradigm

The parameters characterizing high-performance PDs for optical communications are low dark current, high quantum efficiency, and high bandwidth.

Dark current originates from thermionic emission of carriers as well as from bulk, interface, and processing-induced defects. Temperature, band gap, and intrinsic carrier density determine how many carriers are thermally emitted in band-to-band generation as well as generation from mid-gap states; lower band gap materials have higher thermionic emission. Bulk, interface, and processing-induced defects, such as dangling bonds and damage from ion implantation, also act as generation sites that increase dark current. These defects also act as recombination centers that lower quantum efficiency.

PD bandwidth is determined by a combination of transit time and resistive–capacitive (RC) delay limitations. The transit time limited 3 dB bandwidth f_t is given by

$$f_t = \frac{0.45v}{x_d}$$

where v is the saturation electron drift velocity in Ge, 5×10^6 cm/s and x_d is the depletion width or transit distance.

The 3-dB bandwidth in the RC-limited regime is given by

$$f_{\mathrm{RC}} = \frac{1}{2\pi R_L C}$$

where R_L is the load resistance and C is the device capacitance, which is determined by $\varepsilon = 16$, the dielectric constant of Ge; A, the device area; and the depletion width:

$$C = \frac{\varepsilon A}{x_d}$$

Waveguide-integrated PDs typically show $\sim$90% quantum efficiency at near zero reverse bias [3,4], and this high performance is enabled by a core design paradigm of separating the optical absorption path from the carrier collection path. Photon absorption occurs along the longitudinal direction, or length of the detector,

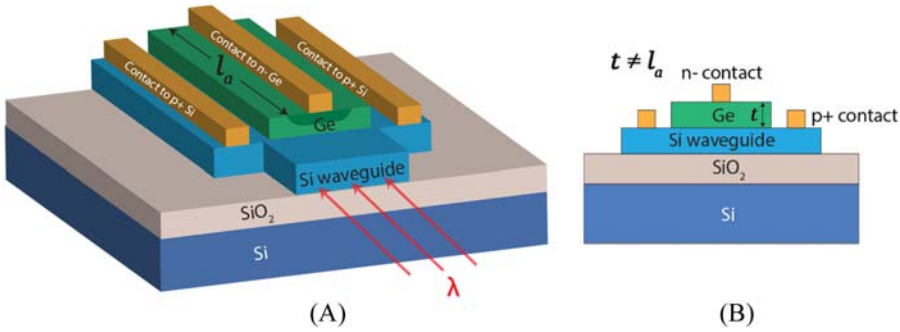

Figure 4.1 Waveguide-integrated p−i−n Ge-on-Si photodetector (evanescently coupled integration scheme) schematic in (A) three-dimensional profile and (B) cross-section. l_α is the effective absorption length and t is the Ge thickness. Typical waveguide integrated photodetector in-plane dimensions are 0.5 μm × 25 μm with a $t = 0.5$ μm Ge thickness. The absorption length l_α is independent of thickness t. The transit time is determined by thickness, and the device capacitance is 5−10 fF.

and photo-generated charged carrier collection occurs in the transverse and/or lateral directions (see Fig. 4.1). The photon absorption is characterized by either an incident butt-coupling scheme, or increasingly by an evanescent-coupling scheme, in which there occurs a gradual transference of a propagating optical mode from the Si waveguide, into an adjacent cladding layer of higher refractive index Ge (see Fig. 4.1). This form of evanescent mode *leakage* from Si waveguide to Ge slab results in an *effective* absorption length for light in Ge, of typically 10 μm.

The optical coupling design is critical to avoid scattering of light at the onset of the detector. Effective refractive index and optical mode matching at the input facet is accomplished by lateral dimension design of the Ge. Furthermore, absorption can be tailored by waveguide mode confinement and consequent extent of evanescence, to ensure the photo-absorbed power does not appreciably contribute to saturation photocurrent—that is, avoid photo-generation of a high enough carrier density such that the charged carriers' electric fields begin to appreciably screen the PD depletion junction field, thereby decreasing carrier collection efficiency.

The schematic for a free space PD device, as used in a fiber optics communications receiver, is shown in Fig. 4.2. Here, the photon absorption occurs along the same direction as photo-generated charged carrier collection. This coincidence fundamentally constrains such a device design with a responsivity versus bandwidth trade-off: the longer the design of the absorption length, the higher the responsivity and lower the 3-dB bandwidth (PD speed limited by both RC delay and carrier transit time).

Returning to the waveguide-integrated detector: here, a more subtle design trade-off exists, as shown in Fig. 4.3. With waveguide integration, the film thickness of a Ge p−i−n PD (see p−i−n structure description in the next section) may be selected to be, at a minimum, comparable to the mode field size of the light signal as it propagates in this Ge slab (at single mode cut-off). While the small thickness size

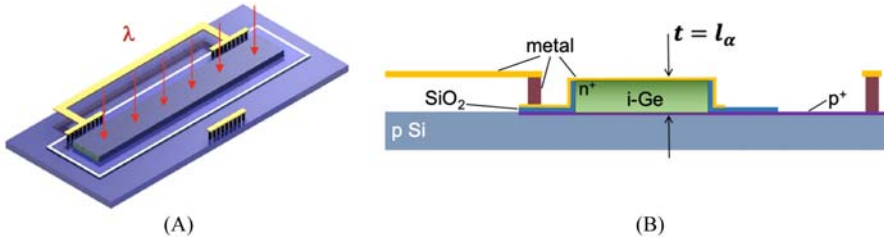

Figure 4.2 A free space Ge-on-Si photodetector, characterized by a larger size, thicker Ge and higher capacitance: photodetector $>10\ \mu m \times 80\ \mu m$, $t > 2\alpha \sim 2.5\ \mu m$, and $C = 50$ fF.

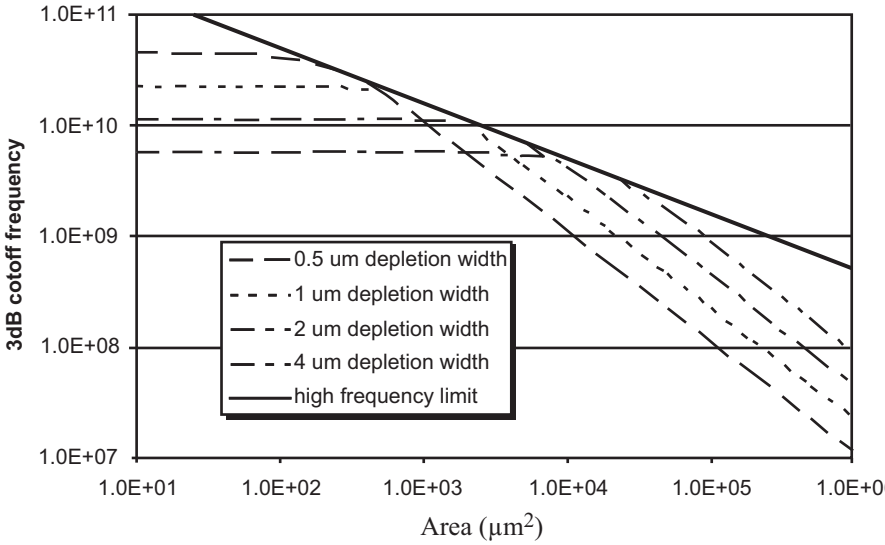

Figure 4.3 Dependence of 3-dB frequency response of Ge p−i−n diodes on device planar area, for various depletion width designs.

provides a small carrier transit time, it results in a larger capacitance from the thin depletion zone, per unit length of the PD. Hence, as the PD length is extended to ensure full absorption of the evanescently coupled optical mode, the device area and capacitance increase, and RC delay begins to exclusively limit the 3-dB bandwidth. The onset of RC delay in dominating the 3-dB bandwidth is observed to occur at a lower frequency, for smaller depletion width p−i−n PDs.

4.1.2 Device structures

A vertical p−i−n photodiode configuration is the most commonly used device structure (see Figs. 4.2 and 4.4A) for high-speed Ge/Si PDs, due to the convenience

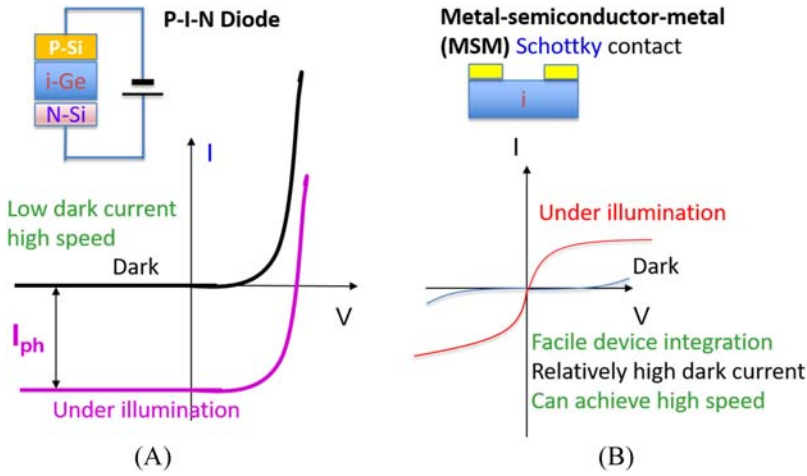

Figure 4.4 Schematic diagrams of device structures and $I-V$ curves of (A) p−i−n photodiodes and (B) metal−semiconductor−metal photodetectors.

of fabrication and uniform electric field in the Ge absorber region. The built-in electric field in the Ge region also enables high-speed operation in photovoltaic mode at 0 bias [5,6]. When integrated with waveguides, a 40-Gb/s data rate has been demonstrated at 0 bias with a responsivity of 0.8 A/W at 1550 nm [2]. A high bandwidth >60 GHz has also been achieved at 1∼2 V reverse bias [6,7]. High-speed Ge/Si lateral p−i−n photodiodes have been investigated by Kotura, Inc. for butt-coupled integration with large cross-section Si waveguides to facilitate coupling to optical fibers. This fabrication requires ion implantation at tilted angles to achieve uniform doping on the sidewalls (see Section 4.3.1 for the description of sidewall processing design, to enhance responsivity and bandwidth) [8,9]. In this case, >36 GHz 3-dB bandwidth has been achieved.

Metal−semiconductor−metal (MSM) is another commonly used Ge/Si PD structure (see Fig. 4.4B). While the dark current of MSM PDs tends to be notably higher than p−i−n diodes due to Fermi level pinning at the metal/Ge interface [3,10], they offer more facile device integration since no high-temperature in situ doping or dopant activation annealing is needed. This simplicity greatly facilitates back-end-of-line (BEOL) device integration as well as strain engineering using SiN_x stressors, deposited laterally about the Ge to create tensile stain and thus increase absorption at a given wavelength [10]. To alleviate the band-bending effect of Fermi level pinning, a thin dielectric layer (∼1 nm) can be inserted between Ge and the metal contact, such as Al_2O_3 deposited by atomic layer deposition [11]. The oxide layer is so thin that it is still electrically conductive via tunneling. Commonly used Schottky contact metals include Al/TiN/Ti [11,12] and Cu/W [13] (top to bottom sequence of deposition). With adequate Schottky contacts, photoconductive gain has also been demonstrated in these devices [10,14]. The demonstrated bandwidth has reached 40 GHz or higher [14].

For p−i−n and MSM PDs, the carrier transit time limit is determined primarily by valence band hole carriers, due to their lower mobility. Traveling wave photodetectors (TWPDs) and unitraveling carrier (UTC) Ge/Si PDs are effective approaches to alleviate the RC limit and carrier transit time limit, respectively. TWPDs are waveguide PDs with electrode configuration that support traveling electrical waves and RF impedance matching with external circuits [15]. Their bandwidths are no longer limited by the RC time constant, thereby allowing longer devices to achieve higher quantum efficiency and higher saturation photocurrent without sacrificing bandwidth.

UTCs, on the other hand, utilize graded doping in the absorber layer as well to create heterojunction band offsets between Ge and Si, to collect photoelectrons while blocking the low-mobility photo-generated holes, thereby keeping them from contributing to the carrier transport process and reducing carrier transit time. These device structures have important applications in microwave photonic links [16,17], where high linearity, high saturation photocurrent, and high 3-dB bandwidth are simultaneously required under high input optical power. These requirements are hard to meet at the same time with regular p−i−n diode structures. This is because higher saturation photocurrent usually requires a larger device footprint to reduce the photo-generated carrier density and minimize electric field screening, and the larger footprint in turn increases p−i−n diode RC delay.

Ge/Si TWPDs with 0.8 A/W responsivity at 1550 nm, 27.5-GHz bandwidth, and a −1-dB compression current of 8 mA (i.e., −1-dB deviation from ideal linear output) have been demonstrated in recent years [18]. p-Ge(absorber)/i-Si(collector)/n-Si waveguide UTC PDs have also been investigated, where the band offset between p-Ge and i-Si blocks the holes and only allows photoelectrons to transport into the i-Si region. They also showed lower capacitance per unit area compared to their p−i−n counterparts. These UTC devices have demonstrated 0.5−0.7 A/W responsivity at 1550 nm, 30−40-GHz bandwidth, and a −1-dB compression current of 1.5−2.1 mA [17].

For high-speed, high-sensitivity photodetection, Ge/Si avalanche photodiodes (APDs) offer unique advantages since they integrate high optical absorption from Ge with low-noise, high-multiplication gain in Si. A separate absorption, charge, and multiplication (SACM) structure is usually adopted. 42-GHz bandwidth has been demonstrated at 1310 nm with a responsivity of 3.6 A/W, corresponding to a gain of 6 [19]. This device will be discussed in detail in Section 4.4.

Fig. 4.5 summarizes the bandwidth × efficiency × gain product and dark current of different types of Ge/Si PDs. We note that, while dark current density is a good indicator of material growth and device passivation quality, in PICs, it is the *absolute* dark current rather than dark current density that governs performance. Therefore smaller device footprint is an effective approach to overcome the increase in dark current density due to crystal defects induced by the large 4% lattice mismatch between Si and Ge.

Waveguide-integrated PDs, which absorb photons along the longitudinal direction to maximize optical absorption while collecting carriers in the transverse direction to minimize carrier transit time, enable not only high efficiency and high

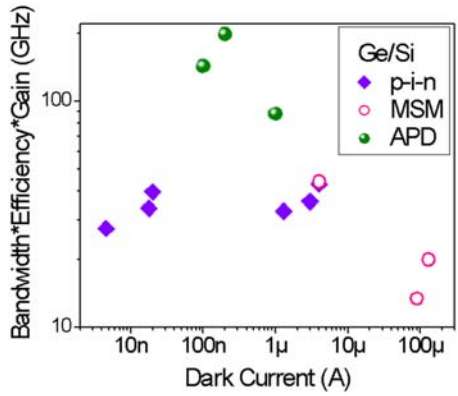

Figure 4.5 Summary of the bandwidth × efficiency × gain product and dark current of different types of Ge/Si photodetectors. The device performance data are collected from Refs [3,19–21].

speed, but also low absolute dark current due to the small device footprint. As shown in Fig. 4.5, the bandwidth × efficiency product of waveguide-integrated Ge/Si p–i–n photodiodes has reached 40 GHz, ~6 × higher than the theoretical limit and >10 × higher than the experimental performance of their normal-incident counterparts, as-summarized in Michel et al. [20]. The dark current of these waveguide integrated p–i–n photodiodes is decreased to 1–10 nA, thanks to the small footprint, compared to μA level for the larger free-space Ge/Si detectors. Therefore waveguide–detector integration realizes both a compatible component for PIC design and leverages the unique benefits of such integration to achieve high responsivity (with high saturation photocurrent), fast bandwidth, and low dark current.

With multiplication gain, state-of-the-art Ge/Si APDs have reached a bandwidth × efficiency × gain product of ~200 GHz, even higher than their III–V counterparts. More detailed comparison with III–V devices will be presented in the next section.

4.1.3 Ge versus III–V high-speed photodetectors

Since InGaA PDs on InP substrate have been widely used in mainstream long-haul telecommunications [22], it is informative to compare the performance of Ge/Si versus InGaAs PDs at optical telecommunications wavelengths. From a material point of view, Ge, an indirect bandgap semiconductor, has two major disadvantages compared to InGaAs. First, its absorption coefficient at 1550 nm is significantly lower. InGaA detectors for optical telecommunications usually have a direct bandgap of ~0.75 eV, ensuring a high optical absorption coefficient at 1550 nm ($h\nu = 0.8$ eV) wavelength light. In contrast, the direct gap of strain-relaxed Ge is exactly 0.8 eV, so the absorption coefficient at 1550 nm is limited by a relatively small joint density of states, close to the bandgap. As a result, the theoretical

bandwidth × efficiency product of vertically illuminated free-space Ge/Si p−i−n diodes is limited to ∼7 GHz at 1550 nm (without backside reflector) [20], compared to ∼20 GHz for free-space (vertically illuminated) InGaAs detectors on InP substrates [23]. Second, the dark current density of Ge photodiodes is significantly higher than that of InGaAs. This is because the indirect band gap of Ge is only 0.66 eV, nearly 0.1 eV smaller than that of InGaAs PDs for telecommunication. Since the dark current density scales exponentially with band gap, this leads to an order of magnitude higher dark current density from Ge photodiodes compared to InGaAs.

Tensile strain engineering and waveguide integration can effectively overcome the above-mentioned disadvantages. For Ge deposited on Si, tensile strain shrinks the Ge direct band gap much faster than the indirect gap, transforming Ge to more closely resemble a direct gap semiconductor [24−26]. Recently Lin et al. [11]. have demonstrated the absorption coefficient of tensile strained Ge can reach or even exceed that of InGaAs in C and L bands, using an SiN_x stressor. Furthermore, as discussed in Section 4.1.1, waveguide integration shrinks the device area and thereby reduces the absolute dark current to 1- to 10-nA level. The separation of the optical absorption path along the longitudinal direction from the carrier collection path in the transverse directions also increases the bandwidth × efficiency product by ∼10× to 40 GHz for Ge p−i−n photodiodes, as shown in Fig. 4.6.

Fig. 4.6 also summarizes the device performance of waveguide-coupled Ge/Si p−i−n photodiodes and APDs compared to their III−V counterparts. Waveguide-coupled Ge/Si p−i−n photodiodes have achieved even higher bandwidth × efficiency product and lower dark current than free-space InGaAs/InP p−i−n detectors [29]. This is a remarkable demonstration of the benefits of waveguide integration in overcoming the limitations of material properties. The performance of waveguide-coupled Ge/Si APDs has also exceeded that of both free space and waveguide-coupled III−V devices. In particular, the low-noise multiplication gain offered by Si is a big benefit for APD performance.

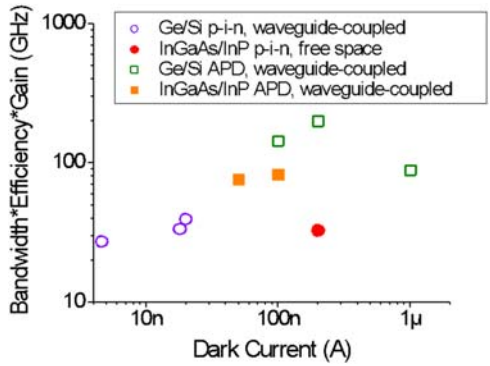

Figure 4.6 Performance comparison of Ge/Si versus InGaAs/InP high-speed p−i−n photodiodes and APDs. The device performance data are collected from Refs. [3,19,22,23] for Ge/Si devices and from Refs. [22,27,28] for III−V devices.

In summary, although the opto-electronic properties of indirect band gap Ge are less ideal than direct band gap InGaAs, tensile strain engineering and waveguide coupling with Si PICs have largely overcome these disadvantages of Ge, leading to device performance comparable to or better than InGaAs/InP PDs for photonic datalinks.

4.1.4 TxRx subsystem requirements

The energy of data transport is becoming larger than the central processing component as systems approach "exascale" FLOP/s and B/s. As an example, $200\times$ more energy is needed to *transport* a bit to a nearest-neighbor chip than is required to *operate* on it. Across a multistage data center network, with routers and multiple transceiver hops, the factor may become $1000\times$. The estimated energy for a FLOP is $0.05-0.1$ pJ/bit. In contrast, the estimated energy for data transport is $2-10$ pJ/bit. Hence, the power budget for a typical electrical link includes $\sim 20-25$ dB loss ($>99\%$) between the transmitter and the receiver. Photonic communication at the board level has thus become an ultra-high-performance computing system design rule. The implementation of intrachip photonic communication will be driven by (1) a saturation of power efficiency scaling for metal at short distances and (2) the WDM bandwidth density advantage of photonics.

Three subsystem factors influence the evolution of chip-level PD integration into data center systems and networks: optical power, electrical power, and parasitic contributions to noise and efficiency. Cost, bandwidth density and power are the primary scaling metrics, and latency is an important functional metric. A hierarchy of latency requirements will determine the locality of data processing. Autonomous vehicles require on-board sensing, training, simulation and actuation; while stored "compression" algorithms can be accessed from the data center. As manufacturing tools become more autonomous, data center control will emerge. For extreme latency constrained functions, RF photonic communication is desirable, and higher optical power than used for digital systems is required to meet noise requirements. For communications link power budgets, the receiver transimpedance amplifier (TIA) dominates power consumption. Electronic–photonic integration can mitigate against capacitive parasitics by using local interconnect design used commonly for transistors. Phototransistors for embedded amplification are large relative to single PD device with associated TIA circuitry. In general, power consumption is proportional to the number of amplification stages in the TIA. The antidote for electrical power consumption is higher optical power and fewer amplification stages.

For ultrafast (>100 GHz), high-power PDs, Ge/Si still lags behind InGaAs/InP in that the bandwidth and the saturation dark current of UTCs and TWPDs are still notably worse than III–V devices for higher power microwave photonics applications. As discussed in the previous section, currently, the bandwidth of Ge/Si TWPDs and UTC is limited to $\sim 30-40$ GHz at a saturation current of $2-8$ mA [17,18], compared to $>$100-GHz bandwidth at $>$10-mA saturation current and

60-GHz bandwidth at >40-mA saturation current from their III−V counterparts on InP [16,22]. That said, Ge/Si PDs allows facile, large-scale integration on a Si platform. Therefore the overall performance could be compensated by using dense Ge/Si detector arrays in this case, benefiting from a much higher thermal conductivity of Si compared to InP [16].

4.2 Epitaxy: dislocations and lattice strain in Ge/Si

4.2.1 Epitaxial growth of Ge on Si by chemical vapor deposition

An epitaxial layer of Ge on Si is typically used for near-infrared p−i−n PDs in integrated Si photonics [29]. The crystalline quality of Ge is the key factor for high photodetection efficiency and low dark current. The Ge thickness is typically several 100 nm in waveguide-coupled PDs, which is much larger than the critical thickness of <10 nm. The critical thickness is defined as the layer thickness of a Ge thin film that is deposited on a Si substrate, at which its volume strain energy, due to the high 4.2% lattice mismatch between Ge and Si, exceeds dislocation line tension energy. At this critical thickness, it is thus energetically more favorable to nucleate dislocations, to relieve the severe lattice mismatch.

A uniform Ge layer can be epitaxially deposited on Si by chemical vapor deposition (CVD) via a low−high temperature two-step growth technique, as Luan et al. [30] reported. In contrast to the thick graded SiGe buffer approach [31], a buffer layer of pure Ge as thin as 50 nm is directly grown on Si at a low temperature of 300°C−400°C, followed by a higher-temperature growth (typically 600°C). The low-temperature buffer layer prevents three-dimensional nucleation of Ge in the Stranski−Krastanov islanding growth mode. Ge can thus be epitaxially deposited via blanket deposition over a whole Si wafer, or via selective epitaxial growth (SEG), where Ge is deposited into local trenches opened up (to reveal the underlying Si substrate) within an intermediate layer of SiO_2.

The SEG process is similar to the Damascene process used for embedding metal structures in dielectrics. SiO_2 is either thermally oxidized or deposited, and then patterned to the ideal Ge dimensions. The SiO_2 is then etched to form a trench, which Ge then fills during growth. Chemo-mechanical polishing (CMP) is then conducted to planarize the surface. Selective epitaxy is desirable since it results in auto-patterned Ge with SiO_2-passivated sidewalls. Fig. 4.7A shows a typical cross-sectional scanning electron microscope (SEM) image of SEG of Ge on an exposed Si (001) surface, in an SiO_2-masked wafer [32,33]. The deposited Ge mesa structure is composed of a (001) plane at the top and a {113} facet plane at the sidewall. Such a Ge SEG mesa structure has been widely used in the fabrication of waveguide-integrated Ge PDs, on Si-on-insulator (SOI) platforms [34−41]. Fig. 4.7B shows a typical top-view optical microscope image of a waveguide-integrated Ge PD along with a cross-sectional schematic [32]. A corresponding cross-sectional transmission electron microscope (TEM) image is shown in Fig. 4.7C [42].

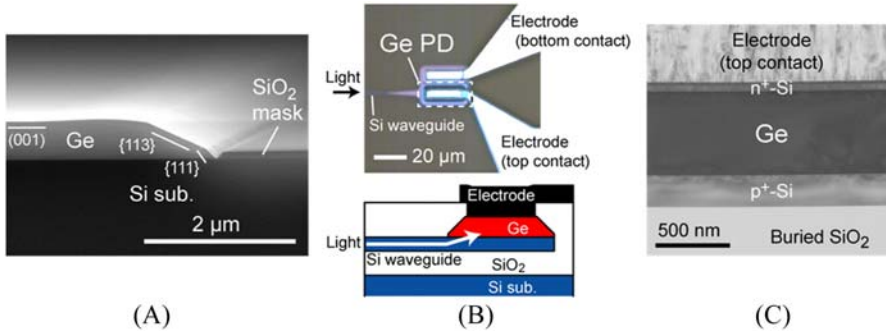

Figure 4.7 (A) A cross-section scanning electron microscope image of Ge selectively deposited on Si within a SiO_2 mask [32,33]. (B) A top-view optical microscope image of a waveguide-integrated photodetector on an Si-on-insulator substrate, and a cross-section schematic [32]. (C) A cross-section transmission electron microscope image of an integrated Ge photodetector on Si [48].

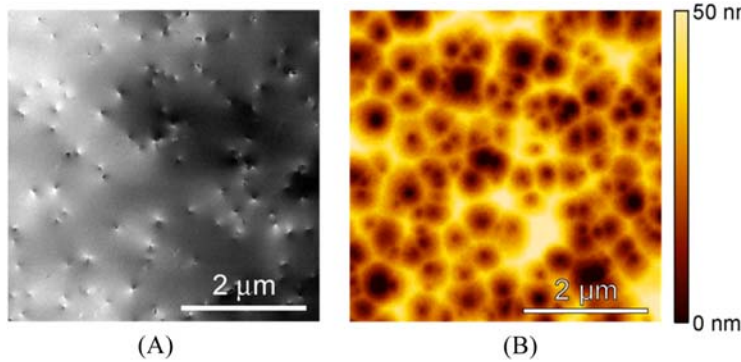

Figure 4.8 (A) A plan-view transmission electron microscope image of as-grown Ge on Si. (B) An atomic force microscope image of as-grown Ge on Si, with etch pits formed after dipping into a $HF/HNO_3/CH_3COOH/I_2$ solution.

4.2.2 Threading dislocation density in Ge

Despite the Ge layer being uniform in thickness, threading dislocations (TDs) are generated, which act unfavorably as generation and recombination centers of carriers. Fig. 4.8A is a typical plan-view TEM image of a 1-μm thick Ge layer, showing a high TD density (TDD), which generate the observed surface defects [42]. Fig. 4.8B shows an atomic force microscope image of the Ge layer after dipping into a $HF/HNO_3/CH_3COOH/I_2$ solution for several seconds. A counting of the defects' resulting etch pits shows a TDD as high as $4 \times 10^8\ cm^{-2}$ [30,43,44]. X-ray diffraction (XRD) is another effective technique for evaluating material quality [45], where TDD is approximately expressed as $TDD \sim \beta^2/(9b^2)$ [46]. Here, β is the

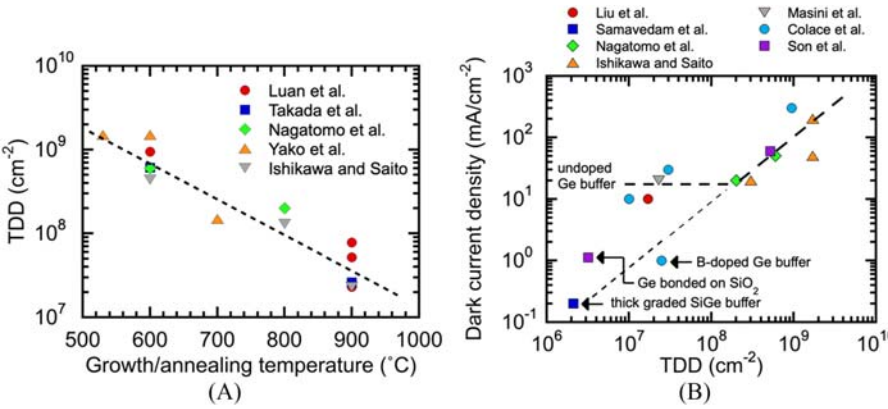

Figure 4.9 (A) Threading dislocation density as a function of the growth/annealing temperature (data plotted from Refs. [30,42–44,48]). (B) Areal dark current density of Ge pin photodetector as a function of threading dislocation density (data plotted from Refs. [26,31,42,44,50–53]).

FWHM in an X-ray rocking curve (ω scan), and b is the absolute value of the Burgers vector (0.40 nm).

TDD tends to decrease with increasing Ge thickness [47,48]. For the several 100 nm of Ge layers used in PDs, postgrowth annealing is usually performed at temperatures of 800°C–900°C (c.f., Ge melting point of 937°C). Fig. 4.9A shows TDD as a function of the growth/annealing temperature. The annealing reduces TDD down to $\sim 10^7$ cm^{-2} level [30,42–44,48]. SEG epitaxy results in lower TDD than blanket epitaxy since dislocations have less distance to glide before annihilating at a surface [30,43]. TDD's impact on dark current, derived by Giovane et al., is the photodiode leakage current per unit of dislocation length:

$$J_{\mathrm{TD}} = \frac{J_0{}^{\mathrm{bulk}}}{(\mathrm{TDD})W} = 2.1\ \mu\mathrm{A/cm}$$

where $J_0{}^{\mathrm{bulk}}$ is the bulk leakage current density and W is the depletion width of the diode [49].

Fig. 4.9B summarizes reports of TDD reduction leading to a significant decrease in dark leakage current for p–i–n PDs [26,31,42,44,50–53]. Careful interface control has reduced the TDD and thus further decreased the dark current density to as low as 1 mA/cm^2 or below [31,52,53].

4.2.3 Tensile lattice strain by thermal stress

The 4.2% compressive strain due to lattice mismatch in Ge thin films deposited on Si substrates is largely relaxed via the generation of misfit dislocations at the initial stage of growth, which occurs at an elevated temperature. Subsequently, a *tensile*

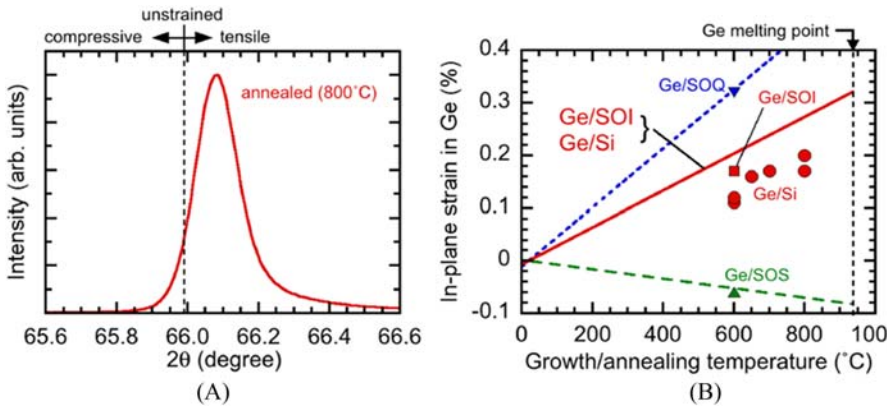

Figure 4.10 (A) X-ray diffraction ω−2θ curve for the 004 diffraction. A Cu Kα radiation was used as the X-ray source in the X-ray diffraction measurements. (B) In-plane lattice strain as a function of growth/annealing temperature.
Source: (A) Reproduced from Y. Ishikawa, K. Wada, Germanium for silicon photonics, Thin Solid Films 518 (6, Suppl. 1) (2010) S83−S87, https://doi.org/10.1016/j.tsf.2009.10.062, with permission from Elsevier. (B) Reproduced from K. Noguchi, M. Nishimura, Y. Tsusaka, J. Matsui, Y. Ishikawa, Enhancement of L-band optical absorption in strained epitaxial Ge on Si-on-quartz wafer: Toward extended Ge photodetectors, J. Appl. Phys. 128 (13) (2020) 133107, https://doi.org/10.1063/5.0023703, with permission from AIP Publishing.

strain is induced into the largely relaxed film, as it cools to room temperature, due to the thermal expansion mismatch between Ge and Si [24,54,55]. The Si substrate has a smaller coefficient of thermal expansion (CTE), $\sim 2.6 \times 10^{-6}\,\mathrm{K}^{-1}$, than the overlying Ge layer, CTE $\sim 5.8 \times 10^{-6}\,\mathrm{K}^{-1}$, and this generates a tensile biaxial stress in the Ge film layer, as the Ge lattice constricts faster than Si upon cooling to room temperature, after growth. Fig. 4.10A shows typical ω−2θ XRD curves for an annealed Ge layer [56]. The Ge (004) diffraction peak is located at a diffraction angle larger than that for unstrained Ge. This indicates a decrease in the out-of-plane lattice constant and an increase in the in-plane lattice constant. As shown in Fig. 4.10B, the in-plane strain tends to increase with the higher growth/annealing temperature. The tensile strain is as large as 0.2% for annealed Ge layers grown on Si and SOI wafers [55−57]. An increase in the CTE mismatch with using a Si-on-quartz wafer (CTE of $\sim 0.5 \times 10^{-6}\,\mathrm{K}^{-1}$ for the underlying quartz substrate) leads to an increase in the Ge tensile strain, while a decrease in the mismatch with using a Si-on-sapphire wafer (CTE of $\sim 7 \times 10^{-6}\,\mathrm{K}^{-1}$ for the underlying sapphire substrate) leads to a decrease in the Ge tensile strain [57].

The tensile strain induces a reduction in the direct band gap energy of Ge. Fig. 4.11A shows a calculation result [56] based on the formulas derived from the k·p method (as-described in Ref. [58]) where the deformation potential values used were $a = -8.97$ eV and $b = -1.88$ eV [25].

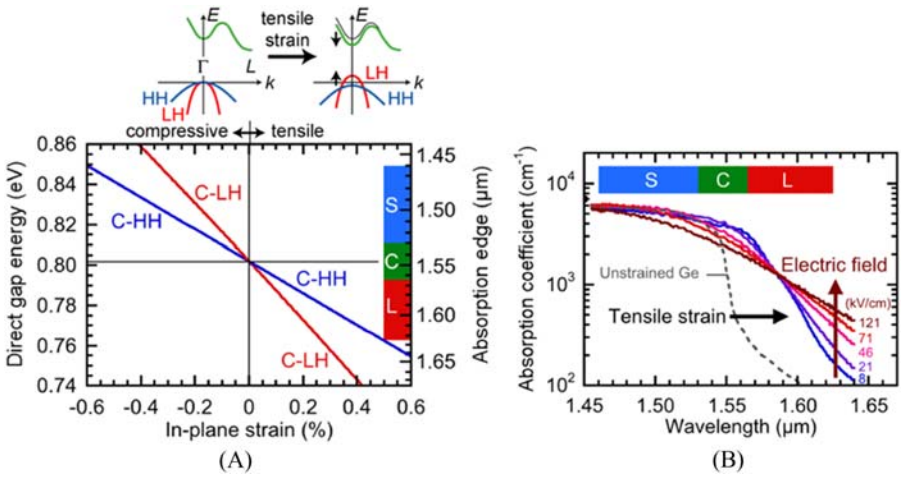

Figure 4.11 (A) Direct bandgap energies between the conduction (*C*) and light/heavy-hole (*LH*/*HH*) valence bands. (B) Optical absorption spectra.
Source: (A) Reproduced from Y. Ishikawa, K. Wada, Germanium for silicon photonics, Thin Solid Films 518 (6, Suppl. 1) (2010) S83−S87, https://doi.org/10.1016/j.tsf.2009.10.062, with permission from Elsevier. (B) Reproduced from J. Fujikata, M. Noguchi, K. Kawashita, et al., High-speed Ge/Si electro-absorption optical modulator in C-band operation wavelengths, Opt. Express 28 (22) (2020) 33123−33134, https://doi.org/10.1364/OE.405447 © The Optical Society.

In the case of 0.2% tensile strain, the direct band gap energy is reduced from 0.80 to 0.77 eV. As seen in Fig. 4.11B, this band gap narrowing causes a shift in the optical absorption edge from 1.55 to ∼1.61 μm, enabling efficient photodetection across the entire C band (1.530−1.565 μm), in addition to shorter wavelength ranges, such as the O band (1.260−1.360 μm). Optical absorption is also enhanced in the L band (1.565−1.625 μm), where the electro-absorption effect, or the Franz−Keldysh effect under high electric field (> 10 kV/cm), also plays an important role [59,60].

4.3 Advanced Ge photodiode processing on CMOS platform

Alongside the question of how to process Ge thin films, another important problem to address is how to design the whole Ge device integration flow. Finding the optimal way to achieve the best structural and processing compatibilities among various devices on a chip and achieving the best overall performances is a complex problem and an important consideration in the field of integrated photonics. Current incumbent integration schemes shown in Fig. 4.12 will be reviewed, in addition to other approaches with potential in future Si photonics integration.

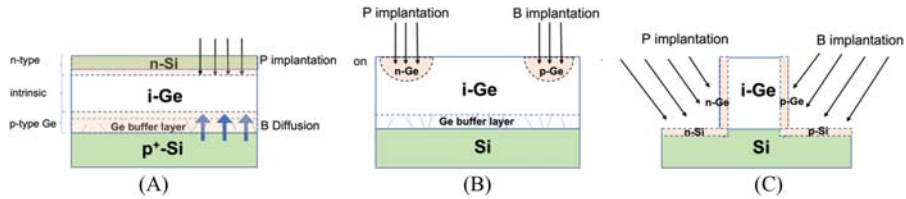

Figure 4.12 Schematic representation of typical doping schemes adopted in many Ge p−i−n photodetectors (A) vertical p−i−n, (B) lateral p−i−n, and (C) lateral p−i−n with side electrodes structure.

4.3.1 Processing challenges with Ge photodetectors

For Ge PDs, a p−i−n photodiode structure is most typical. As a doping method to form p−i−n junctions in Ge film, modulating PH_3 or B_2H_6 dopant gas flow for in situ doping and junction forming during the Ge CVD growth is not desirable due to autodoping from the reaction chamber [61]. Instead, an intrinsic Ge epitaxy followed by ex situ doping process can be more practical, especially if the various types of Ge photonic devices (e.g., PDs and lasers) with device-specific doping profile designs are to be cointegrated on one chip.

Fig. 4.12A shows the typical Ge film doping structure and process that are widely adopted for many vertical Ge p−i−n photodiodes. Normally, the epitaxial intrinsic Ge is deposited on a highly doped Si wafer substrate [5], or on the upper Si layer of an SOI wafer which has been implanted with dopants beforehand [40,62−65]. It is hypothesized the lower region of the Ge film would be doped by dopant diffusion from the substrate, during the high-temperature Ge epitaxial growth. The upper region of the Ge film is typically doped by ion implantation, and subsequently activated by annealing.

Most demonstrated vertical p−i−n PDs have adopted a p-type electrode at the bottom and a n-type electrode at the top. It is important to note the diffusion of boron (B) in Ge has been recently reported to be significantly slower than in Si, due to its interstitial-mediated diffusion mechanism; [66,67] while the diffusivities of donors, such as phosphorus (P) and arsenic (As), in Ge are enhanced relative to those in Si. The reported measurements were done in bulk Ge and Si wafers. Hence, it may be expected the actual B diffusion in a Ge buffer layer, deposited on Si, will be relatively faster due to the high concentration of misfit dislocations, which will provide high-speed diffusion pipe paths. A quantitative experimental verification of such enhanced diffusion in the low-temperature-grown Ge buffer layer has yet to be done, to our knowledge.

If the dopant diffusion from the substrate is enough, the lower portion of the Ge film may be fully doped to form an electrode, thereby reducing the negative impact of the high defect density region on PD performance (i.e., generation of minority carriers in the intrinsic depletion region due to the Shockley−Read−Hall process, resulting in dark current). The formation of high doping concentration, shallow n-type Ge with low-resistivity contact to metal is a significant challenge [68,69].

Inserting a n^+-Si interlayer between the top Ge layer and metal (see Fig. 4.12A) has been widely adopted as an effective option in many devices [5,62–64]. An alternative is to use an n^+ Si substrate and form a vertical Ge p–i–n PD with n-type on the bottom and p-type on the top [65], in which case the direct contact of $p+-Ge$ with metal can be done conveniently without an Si interlayer.

A lateral p–i–n PD structure is another popular scheme [41,70]. With this type of arrangement, p-type and n-type electrodes can be formed side by side on the same top surface of Ge via ion implantation and activation annealing, as shown in Fig. 4.12B. Having metal contacts and electrodes at the same height level may provide processing simplicity and structural convenience, especially in large-scale integration circuit settings. Whereas a vertical p–i–n PD structure as in Fig. 4.12A may have Si–Ge heterojunctions that depend on the diffusion length of p-/n-type dopants from Si into the Ge layer, the lateral structure will have Ge p–i–n homojunctions, which is claimed to be relatively advantageous in terms of responsivity and speed [41].

On the other hand, the existence of an undoped Ge buffer layer region with high dislocation density in the depletion region may cause significantly higher dark current. Furthermore, if the p-type/n-type electrodes are formed on the top surface, electrical field line penetration may be weaker in the lower part of the Ge layer, lowering the photocarrier collection speed from this region. It could be a viable option to conduct deep dopant ion implantation to form deep vertical electrodes or use a narrow Ge strip structure combined with side-electrode formation [70] as shown in Fig. 4.12C. However, a side implantation device design involves more processing complexity.

Studies on dopant diffusion in Ge have focused mostly on controlling diffusion to the smallest scale. This is especially the case in the area of Ge MOSFET transistor research, because it is critical to form ultra-shallow junctions for nanoscale devices. The doping process for a Ge MOSFET is predominantly done by ion implantation. In contrast, Ge photonic devices often have larger device dimensions (in the order of the wavelength divided by refractive index, for example, $\sim$400 nm) and can generally be less restrictive to diffusion. Therefore PD processing may take advantage of drive-in dopant diffusion.

For example, diffusing dopants from highly doped Si electrode layers, deeper into the underlying Ge, results in Ge p–i–n homojunction PDs with Si electrodes terminals for contact to metal. The benefit of diffusion doping over implantation is that it does not cause any subtle lattice damage or defects in Ge, reducing detector leakage current. Photoluminescence intensity and lifetime measurements can indirectly assess the effect of high-concentration diffusion doping, while preserving the Ge lattice: any lattice damage and defects in Ge will serve as nonradiative recombination centers that reduce the luminescence lifetime. For example, phosphorus diffusion from a phosphosilicate glass source covering the Ge film [71] demonstrated a drastic PL intensity enhancement of as much as 10$\times$ (greater than any values reported), versus an intrinsic Ge reference sample (see Fig. 4.13A). In contrast, phosphorus doping by implantation and activation annealing may have some concerns of limited or relatively moderate photoluminescence enhancement [72–74].

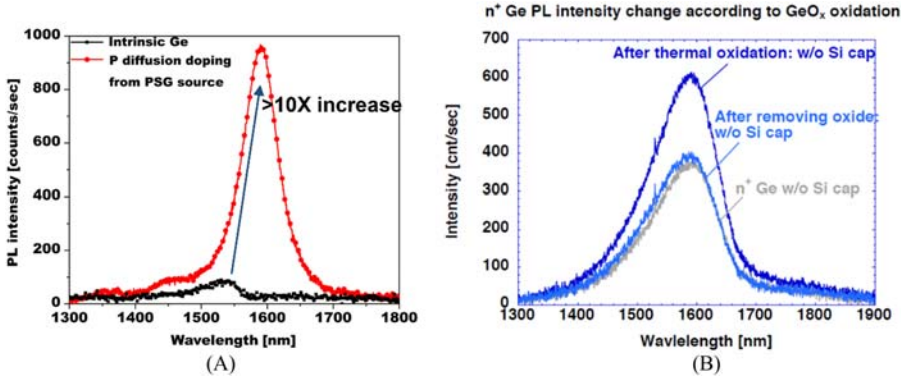

Figure 4.13 (A) Photoluminescence (*PL*) intensities of Ge-on-Si film after phosphorus diffusion doping process versus undoped as-grown Ge sample [71]. (B) PL intensity of the n + Ge with a surface as-grown, after thermal oxidation, or after removing the oxide [75].

In addition to the efforts to minimize the lattice damage and defect density in Ge, much attention needs to be paid to the surface passivation of the Ge PD. Several dark current analysis studies of Ge p−i−n PDs [53,65] indicate that in addition to bulk leakage current (which is directly associated with Ge film epitaxial quality), surface leakage current is a major or even larger contributor to the total dark current. This suggests effective Ge surface passivation is key to achieving the improved PD performance, especially when the Ge film quality improves (i.e., TDD is below a threshold level). It has been shown that Ge surface passivation treatments, such as postmetallization annealing [76,77] or intentional surface doping [77], can significantly reduce the dark current of a Ge PD. Recently Pan et al. [75] showed the greatest enhancement of photoluminescence from the shallow surface region of a Ge film after it was thermally oxidized to form a GeO_x surface (see Fig. 4.13B), suggesting another potential path towards optimal Ge surface passivation [78]. Such advanced processing for Ge surface passivation may play an important role in pushing Ge PD performance to its fullest potential in integrated photonics.

4.3.2 Advanced integration of Ge devices with CMOS electronics

The vast majority of demonstrated electronic and photonic cointegrated circuits (EPIC) that have adopted Ge PDs were implemented on SOI wafers. SOI wafers provide an ideal platform for PICs due to the high crystalline quality of the top Si layer, which can serve as the core medium for a Si waveguide whose optical undercladding is the buried oxide (BOX) layer. Waveguide-integrated Ge PDs can be grown and fabricated on top of such SOI-based Si waveguides. At the same time, CMOS electronic circuits can be fabricated on SOI wafers using standard Si processing. Using this wafer platform, fully monolithic EPIC transceivers on a CMOS platform have been demonstrated [79].

However, even with SOI wafers, there is still some mismatch between Si electronics and Si-compatible photonics, in terms of their preferred substrate, structure, and processing. First of all, nearly all mainstream digital ULSI electronic chips are based on the Si bulk wafer. Even if electronic circuits are to be manufactured on SOI, SOI wafers that can provide benefits, such as latch-up prevention and less parasitic capacitance, typically should have an ultra-thin top Si and BOX, both to within a few tens of nanometers. In contrast, photonics SOI has typical specifications of ~200-nm-thick Si top layer with up to 2-μm BOX [80].

This wafer platform mismatch necessitates differentiated approaches to monolithic EPIC integration. If a fully monolithic EPIC chip is to be developed as a photonics-oriented product, for example, an optical transceiver chip for an active optical cable, an SOI wafer with photonics-oriented specification will be used as the integration platform; the electronic circuits formed on this platform will only serve a secondary role, such as driving optoelectronic devices. An optical receiver with extensive waveguide circuitry and diverse kinds of photonic devices, such as spot-size converter (SSC) and demultiplexer (DMUX), integrated on one chip would certainly fall into this category. In fact, many examples of SSC [81] and DMUX [82] devices have been recently developed on photonics-specification SOI wafer platform.

In contrast, if electronic and photonic cointegration is used to enhance the performance of digital ULSI processing chips, with added photonic functionalities, such as optical interconnection and optical I/O pins (i.e., an electronics-oriented product), the electronic-integrated circuit architecture for Si bulk wafer will remain unchanged and photonic devices including active Ge PDs will have to be adapted for better acceptance to CMOS electronics (see Fig. 4.14). Optical interface for processor-to-memory interconnection, as suggested by Samsung [83], is one such example.

One approach to realize this EPIC is to include photonic devices in standard Si CMOS process front-end-of-line (FEOL) integration (see Fig. 4.14A), which is to embed photonic devices on the front-end substrate level on a par with CMOS devices. This approach requires Ge film growth and device processing on a Si wafer surface at a high temperature. With the photonics-oriented SOI platform taken away, the PIC will need to rely on (1) deposited dielectric waveguides and vertical waveguide integration in the BEOL or (2) local modification of the substrate. An example of local modification of the substrate is to fabricate passive photonic devices on top of oxide trenches (shallow trench isolation) [84]. A more aggressive approach is to optically isolate photonic devices from the underlying Si of the electronics-oriented SOI wafer, by locally removing Si from beneath the portion of the ultra-thin BOX layer atop which rest the devices [85]. There is also a more radical approach of photonics integration that requires "zero-change" to standard CMOS electronics processing, by significantly adapting Si photonics and Ge photonic devices into a whole new class of devices that are strictly based on the materials already used in electronics processing [86,87].

Another alternative approach with good future potential is a BEOL integration scheme (see Fig. 4.14B), where a photonics-dedicated plane is formed at an upper

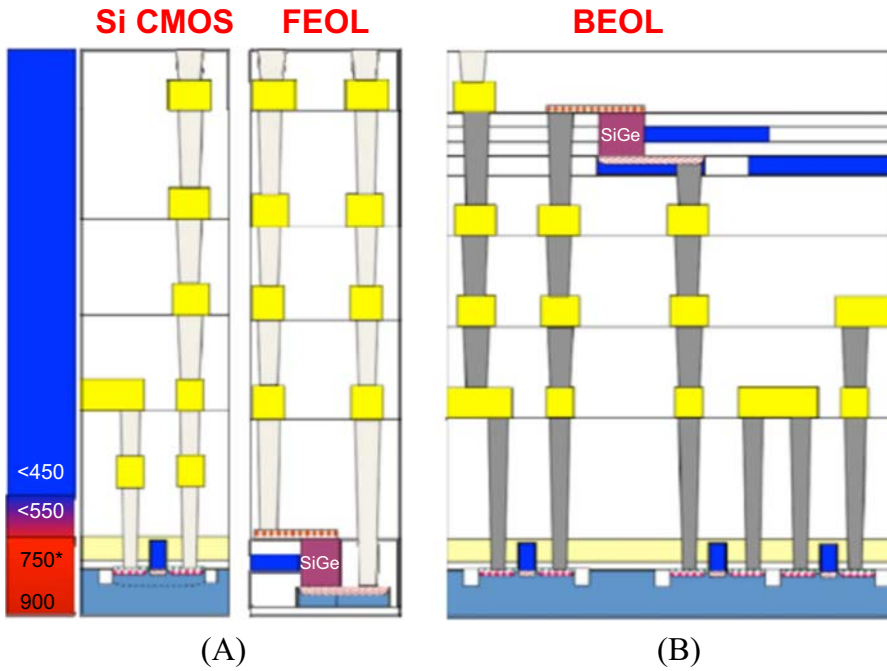

Figure 4.14 Schematics of different schemes for photonic device integration with silicon CMOS process. Insertion of active and passive photonic devices into (A) front-end-of-line plane and (B) back-end-of-line plane.

level of the chip, over the interlayer dielectric (ILD) oxide material. BEOL photonics integration architecture can be, if suitably developed, most compatible with mainstream incumbent ULSI CMOS electronics, because it can alleviate constraints of the FEOL approach (such as real-estate interference at the Si surface between the electronics and photonic routing network), structural incompatibility, and the processing complexity [88]. A significant challenge to this approach is the difficulty in fabricating high-performance photonic devices, such as Ge PDs and high-index-contrast Si waveguides, at the BEOL level [89], due to lack of a crystalline substrate and restrictions on processing temperature. BEOL photonics integration calls for further innovative research on materials processing, an example of which includes single-crystal Ge growth on top of ILD oxide [90].

Other possible approaches that deserve attention may be the utilization of three-dimensional chip bonding between electronics chips and photonics-dedicated chips [91–93] and a FEOL–BEOL hybrid integration architecture that places active photonic devices in the lower substrate level, to leverage a crystalline Si substrate and high processing temperature, while placing passive devices and waveguides on an upper level to avoid interference with the substrate or surrounding electronic components [94,95].

The CMOS electronics' demand for high-bandwidth optical interconnection for intra- and interchip communication is evident, and high-performance Si-/Ge-compatible photonic devices are well developed so far. However, as none of the above-mentioned cointegration schemes of ULSI electronics and photonics are well established yet, there is still a grand challenge of research and development ahead in this area.

4.4 Avalanche photodetectors

4.4.1 Breakthrough and current market: Ge/Si heterojunction avalanche photodiodes in home and office communications

APDs are widely used for fiber optics telecommunications and optical sensing applications. With the development of optical networks, such as Fiber to the Building/Home (FTTx) [96], and sensor markets, such as LiDAR [97], billions of APD devices have been deployed in the past several years. While we will focus on the telecom application, the sensor-related markets are presently a booming revenue industry for APD.

Traditionally, the majority of APD receivers are made of InP-based devices, which provide a 5- to 10-dB sensitivity improvement, compared to p−i−n PD receivers when the operating data rate is $<$10 Gb/s [98]. However, due to the exponentially increasing demand of applications, such as HDTV, AR, VR, and 5-G mobile, the bandwidth requirements of fiber communications today are reaching 400 Gb/s and will grow even further to 800 Gb/s in the next few years [99]. Several new developments in devices and components are being explored, among which the Ge on Si APD (Ge/Si APD) has attracted tremendous attention due to its theoretically higher performance than the traditional III−V APD [100]. The Ge/Si APD has a larger gain $\times$ bandwidth product (GBP) and lower ionization coefficient ratio of holes and electrons (see next section for further details). In addition, the CMOS compatibility of Ge/Si devices provides the possibility of cost-effective solutions for high-speed applications in the telecom market and the promise of mass manufacturing with high performance uniformity within a wafer batch as well as batch to batch. The main challenge of Ge/Si optical devices at CMOS foundries is the development of a low defect density monolithic integration process for pure Ge thin film grown on Si.

In 2015 a 10-Gb/s APD was developed with 1E-12 sensitivity, corresponding to −29.5 dBm at 1550-nm carrier signal wavelength [101]. In 2016 a 25-Gb/s APD was developed with a world record 1E-12 sensitivity, corresponding to −23.5 dBm at 1310-nm carrier wavelength [102]. Lastly, in 2017 a linear 28 Gbaud APD receiver was developed, providing 5- and 5.7-dB improvement over p−i−n PD options for Multi-Mode Fiber and Single-Mode Fiber PAM4 Pulse Amplitude Modulation applications, respectively [103,104]. Even with a large defect density due to the Ge/Si lattice mismatch, Ge/Si APD devices demonstrated good reliability and have passed all required qualification tests according to telecom GR-468

Generic Reliability Assurance Requirements (2004) [105] and have already been widely deployed commercially. A series of APD chips in singlet and array forms and their corresponding Receiver Optical Sub-Assembly (ROSA), with speeds from 10 to 400 Gb/s, have been developed to meet strict market requirements, achieving deployment of more than several million units.

4.4.2 Avalanche photodiode principle and device performance of Ge/Si avalanche photodiode

As mentioned previously, light absorption in Ge and carrier multiplication in Si allows low excess noise and high-speed performance, compared to III−V APDs. APDs operate in high reverse bias to supply sufficient energy to generated photocarriers such that cascades of electron−hole pairs are created by impact ionization. In this way, one photon can generate multiple carrier pairs, thereby generating high levels of electrical signal for low incident photon flux.

The ionization coefficient describes the ability of electrons and holes to impact ionize, and the ionization ratio provides the ratio of the coefficients for hole to electron carriers. It is desirable to have a ratio far from 1, so that one carriers' impact ionization coefficient is much larger than the other, as this results in low noise and high bandwidth APDs.

The basic Ge/Si APD structure has a separate absorption charge and multiplication design (SACM), and it combines the merits of both Ge and Si. The Ge layer can cover optical absorption from 850 to 1600 nm, while Si contributes outstanding avalanche characteristics, with a low ionization coefficient ratio of holes and electrons that results in low excess noise and high GBP. For high-speed and high-performance applications, appropriate internal electrical field distribution design is the key factor to ensure short carrier drift time, low dark current, optimized gain value, and short RC time constant. A wide working bias range is also a key specification that needs to be carefully considered to provide enough engineering margin in actual applications. Fig. 4.15A shows a schematic structure of an SACM Ge/Si

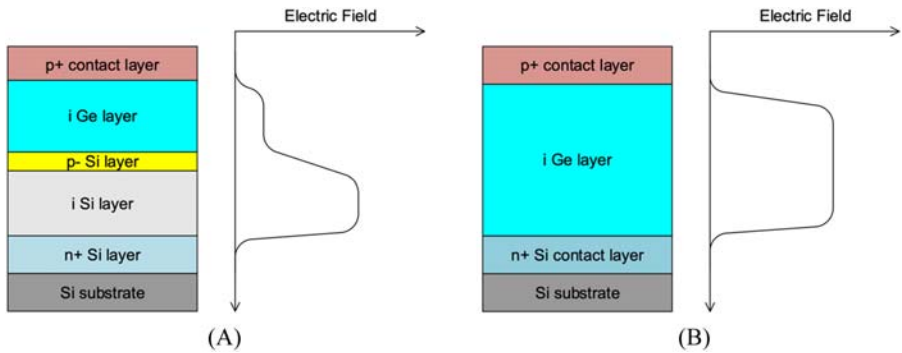

Figure 4.15 Schematic diagram of (A) SCAM Ge/Si APD and its electric field distribution and (B) Ge APD on Si and its electrical field distribution.

APD and its electric field distribution. During operation, a higher electrical field is typically applied in the i-Si layer to trigger the avalanching effect. The Ge/Si APD is distinct from the Ge APD on Si, which relies on Ge for both photo-generation and amplification of carriers (see Fig. 4.15B). Ge/Si APD has intrinsic better performance compared with Ge APD, due to silicon having very low k factor (<0.05) compared with that of Ge ($0.7\sim1$). k factor is defined as the ratio of impact-ionization coefficient, and Si's low k factor is why Ge/Si APD has higher bandwidth and lower avalanche noise compared with Ge APD. And another benefit of Ge/Si APD is the Ge layer can be designed to be biased under very low voltage to realize low dark current operation.

The Ge/Si APD schematic shown in Fig. 4.16A includes in its design several novel features to improve performance: (1) an n-Si layer is inserted between a heavily doped bottom contact layer and intrinsic Si layers, to realize better electrical field distribution, and (2) a backside metal reflector is deposited to enhance responsivity in the target wavelength band. Fig. 4.16B shows the cross-section SEM image of such an APD device.

Fig. 4.17 is the optical microscopic image of 10-, 25-, and 56-Gb/s Ge/Si APD chips, which have an optical incident aperture varying from 20 to 12 μm, all with a backside reflector design. The 10- and 25-Gb/s APD chips were commercially deployed in telecom and datacom systems with application speeds of 10, 40, 25, 50, and 100 Gb/s using Non-Return-to-Zero (NRZ) or PAM4 schemes in a singlet or four-channel array design. Very good performance and high reliability were shown in commercial deployment compared to traditional III−V counterparts, leading to a total quantity of chip and ROSA at more than 2 million, presently.

Fig. 4.18 presents performance plots for the newly developed 56-Gb/s APD chip, including dark current, photo current, and responsivity at 1310 nm, all measured under 25°C. The breakdown voltage (V_{BR}) of the Ge/Si APD is defined as the voltage applied when dark current is equal to 100 μA at 25°C, and the typical value of V_{BR} is ~ -20.5 V. The APD dark current is ~0.7 μA.

Fig. 4.18C shows the RF response of several APD chips packaged in a Transistor-Outline-can (TO-can) semiconductor package type ROSA, with

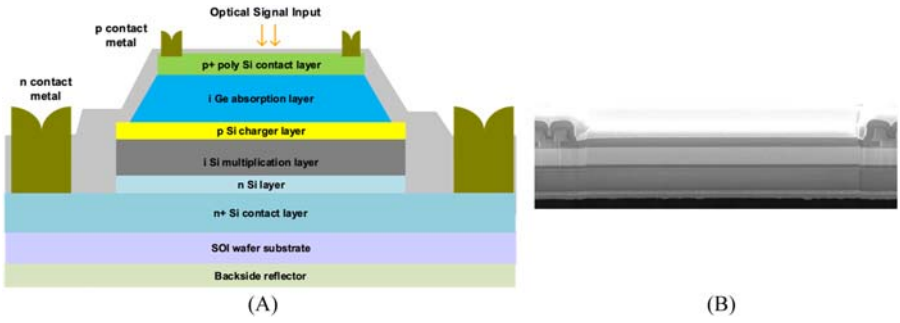

Figure 4.16 (A) Schematic cross-section of Ge/Si APD. (B) Scanning electron microscope image of a Ge/Si APD.

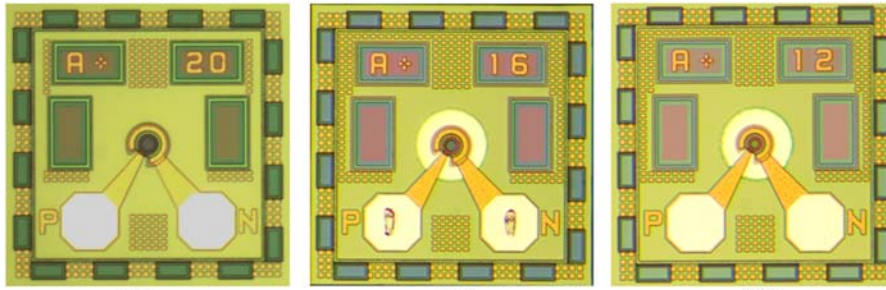

(A) (B) (C)

Figure 4.17 Optical microscopy image of 10-, 25-, and 56-Gb/s Ge/Si APD chip, all realized commercially deployment.

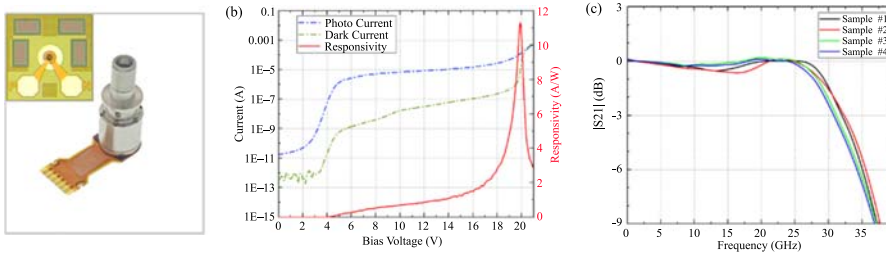

(A) (B) (C)

Figure 4.18 DC/RF measurements of the APD chip and ROSA. (A) Micrograph (inset) and image of a packaged APD chip and ROSA. (B) Photo current, dark current, and responsivity curves of the APD ROSA at 25°C. (C) RF response of the APD ROSA [106].

commercially available TIAs. The 3-dB bandwidths are $\sim$32 GHz at $V_{BR} = -2$ V bias and represent the combination of response limitations of the APD chip, TIA, TO header, flexible printed circuit, and evaluation board.

A 53-Gb/s NRZ sensitivity testing setup is shown in Fig. 4.19A; the inset diagram shows the optical eye diagram of the light source, and the outer extinction ratio is 8.1 dB.

Fig. 4.19B shows the measured back-to-back 53-Gb/s NRZ sensitivity. As reference, a PD ROSA with the same package design was measured first and showed -17.5-dBm sensitivity at 1E-3 bit error ratio (BER). The sensitivity of the APD ROSA was measured to be -23.2 dBm under the same testing setup and BER condition, resulting in a 5.7-dB sensitivity gain compared with the PD ROSA. The inset in Fig. 4.19B is an electrical eye diagram under -10.5 dBm. In 100-Gb/s PAM4 application, APD ROSA also shows similar sensitivity gain compared with PD ROSA, and typical sensitivity of APD ROSA is about -16.0 dBm at 2E-4 BER. A pilot run evaluation including more than 100 pieces APD ROSA shows that it also has good temperature performance and good performance uniformity.

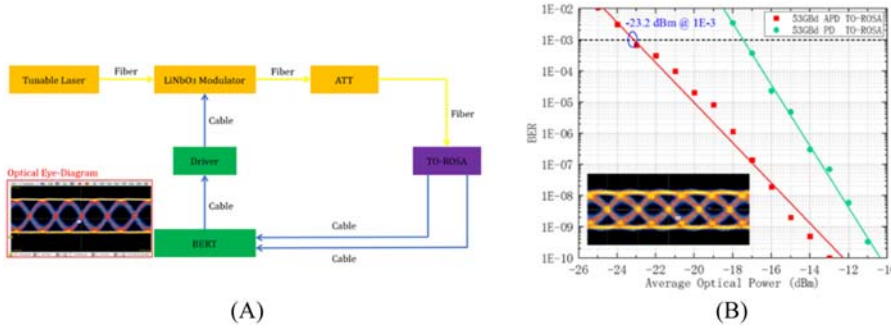

Figure 4.19 53-Gb/s NRZ sensitivity testing of the ROSA. (A) Schematic representation of test setup; inset is the optical eye diagram of the light source. (B) Bit error ratio (BER) measurements in a back-to-back configuration [107].

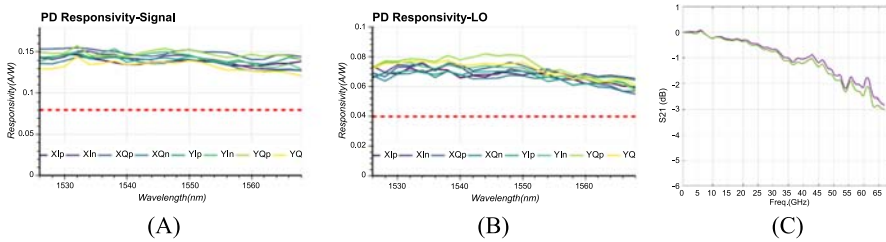

Figure 4.20 64-G ICR PIC responsivity for (A) signal and (B) local oscillationports. (C) Bandwidth of 128-G waveguide-coupled photodetector.

All photonic integrated chip(PIC) based on SOI wafer that include a Tx/Rx transceiver for a photonic data link will require the monolithic integration of a waveguide-integrated p−i−n PD or APD. The material layers' structure of waveguide-coupled APDs is similar to Fig. 4.16A, except that the silicon n + contact layer will not only act as the electrical current flow channel but also act as the optical signal transmission channel of the waveguide. The main structural difference of waveguide-coupled PDs versus APDs is the presence of the silicon multiplication layer. High-speed waveguide-coupled APDs with high performance have been developed. These devices can be monolithically integrated in a DR4 transceiver (400-G 4-channel MZ interferometer chip for $<$2-km link), Integrated Coherent Receiver (ICR) or Integrated Coherent Transmitter and Receiver (ICTR) PIC with nanowire Si waveguides.

Waveguide based detector chips with 32/64/96/128-G Baud rate are now available and have achieved large scale commercial deployment in fiber systems. In addition to the telecommunications market, waveguide-coupled Ge/Si PDs or APDs are attractive devices for sensor applications, such as LiDAR, 3D imaging, and health monitoring (especially in the O and C bands). Fig. 4.20 shows the responsivity of a commercially deployed 64-G ICR PIC and a 128-G compatible waveguide-

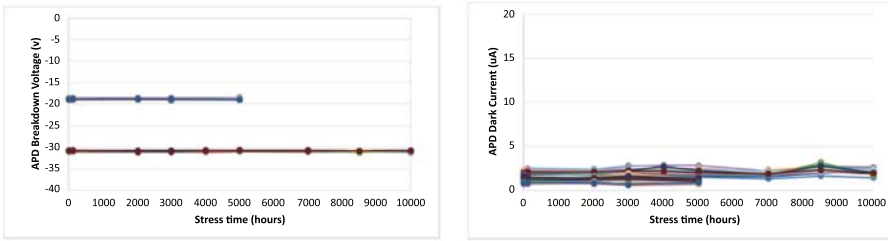

(A) (B)

Figure 4.21 V_{br} and dark current drift plot of 109 pieces of Ge/Si APD subjected to 175°C operations.

coupled PD bandwidth. More than 70-GHz bandwidth (see Fig. 4.20C) for the waveguide-coupled PD implies that it can be used for 112- or 128-G Baud rate applications; this enables up to 200 Gb/s per lane if using PAM4 modulation signals, or 800 Gb/s per lane if using Dual Polarization Quadrature Amplitude Modulation (DP-16QAM) signals. Similar waveguide-coupled APDs in monolithic PIC chips for Intensity Modulation/Direct Detection, coherent communications, or LiDAR applications are anticipated to be released in the coming years.

As discussed in Section 4.2, dislocations formed from the lattice mismatch between Ge and Si result in relatively high dark current for normal incidence Ge/Si PDs; the same holds true for normal incidence Ge/Si APDs. Ge/Si APD dark currents are $\sim\mu$A, in contrast to III−V APD dark currents, which are $\sim$10 nA. However, this higher dark current does not result in sensitivity degradation for the Ge/Si APD because the total noise current of a receiver includes that of both the APD and TIA. The TIA actually contributes the dominant part of noise current, especially for applications >25 Gb/s.

To evaluate the reliability of Ge/Si APDs and estimate their failure rate and lifetime, a high temperature operation test at 175°C is conducted. A total of 109 pieces (pcs) of APD chips from four different wafers have been used in the reliability test, including 10- and 25-Gb/s chips. Fig. 4.21 shows the room temperature breakdown voltage and dark current drift plots of the109 pcs (4 batches) of Ge/Si APDs stressed at 175°C for 10,000 hours. No failure has been observed.

For all Ge/Si APDs mentioned above, besides their high performance metrics, they demonstrated good device reliability by passing all qualification tests required by GR-468 reliability standards [105]. Such high-performing and reliable devices pave the road for current 100-Gb/s and future 400-Gb/s, 800-Gb/s, or 1-Tb/s applications.

4.4.3 Next-generation applications of avalanche photodiode

III−V APDs have demonstrated their significant systems performance in GaAs and InP-based material platforms, as a viable candidate to Ge or Ge/Si APDs. A large GBP of several hundred GHz has been reported in InAlA APDs implemented in

long haul optical telecommunication systems [108]. To integrate III−V APDs on an Si PIC platform, several challenges in growth and processing need to be addressed. One such challenge is the formation of antiphase domains, since III−V epitaxial growth must be done on exact (100) Si. Recently it has been reported that an antiphase domain-free GaA-based APD was prototyped on a saw-tooth structured Si (100) wafer surface [109]. The remaining challenge would be implementation of III−V growth in a Si-CMOS platform.

Ge/Si heterojunction APDs meet current system requirements, as described in Section 4.1.4. The next device milestone for APD design is operation at lower voltage. Although the Ge APD would be an optimal candidate (fewer misfit dislocations without Ge/Si interface within the APD structure), its ionization coefficients of electrons and holes are similar, and this results in relatively more noise after signal amplification. A viable option may be a Ge APD with Ge/Ge_xSi_{1-x} multilayers as the multiplication layer. It has been reported the ratio of ionization coefficients can be engineered using a multilayer structure, by leveraging instantaneous acceleration of carriers at conduction and/or valence band discontinuities (ΔE_C and/or ΔE_V) [110]. For example, ΔE_V is large while ΔE_C is small at the electronic interfaces in a Ge/Ge_xSi_{1-x} multilayer. Accordingly, under an electric field application, holes injected from the Ge_xSi_{1-x} layer will be instantaneously accelerated at these electronic interfaces within the multilayer [111].

Fig. 4.22 shows the ratio of ionization coefficients for electrons to hole, $\frac{\beta}{\alpha}$, increases to ~ 10 when holes are injected into Ge from a Ge_xSi_{1-x} multilayer with $\Delta E_V = 0.2$ eV, under an E-field of 140 kV/cm. When compared with $\frac{\beta}{\alpha} < 2$ in bulk Ge, that is, $\Delta E_V = 0$ eV, electron multiplication in a Ge/Ge_xSi_{1-x} multilayer is

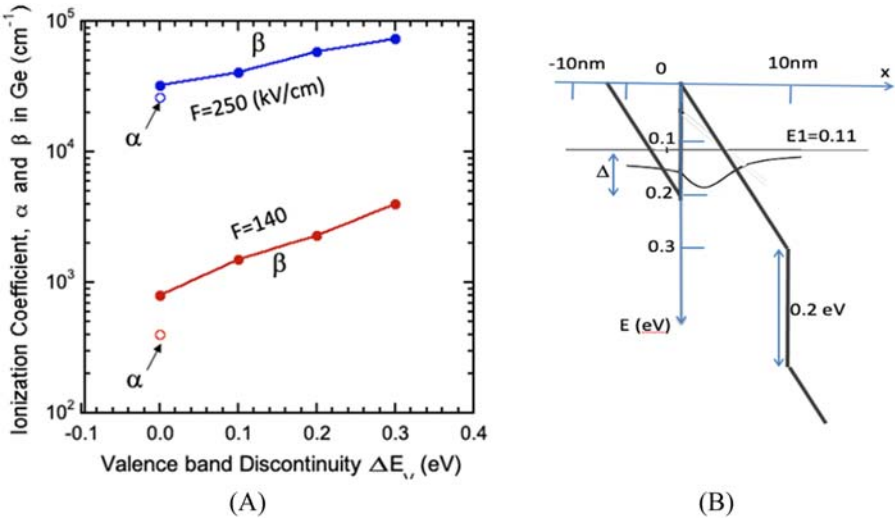

Figure 4.22 (A) Ionization coefficient of holes α and electrons β versus ΔE_V, in Ge. (B) Valence band structure of a Ge/$Ge_{0.4}Si_{0.6}$ multilayer ($\Delta E_V = 0.2$ eV) under a 300 kV/cm E-field.

significantly suppressed. The valence band structure of $Ge/Ge_{0.4}Si_{0.6}$ is shown in Fig. 4.22B. Holes injected from the $Ge_{0.4}Si_{0.6}$ layer are accelerated by $\Delta E_V = 0.2$ eV and are multiplied in the Ge layer. Hole tunneling from Ge to $Ge_{0.4}Si_{0.6}$ is also enhanced by a triangle well, as shown in Fig. 4.22B. The ground subband is located at -0.1 eV in the triangle well, because a hole can tunnel through the barrier and reach the valence band of the Ge_xSi_{1-x} layer. Thus the hole multiplication process is not hindered by the potential barrier. $Ge/Ge_{0.4}Si_{0.6}$ multilayer APDs can multiply holes by $10\times$ at a typical CMOS-compatible applied voltage of ~ 3 V. The Ge and $Ge_{0.4}Si_{0.6}$ layers are both 10-nm thick in this figure, which is less than the critical thickness (i.e., avoid formation of misfit dislocations) of $Ge/Ge_{0.4}Si_{0.6}$ multilayer structure.

4.5 Photodetectors in data centers

The data center is part of a continuum of computer architectures that range from early punch card installations to minicomputers, desktops, laptops, enterprise data centers, and the cloud. The drive towards higher efficiency and special purpose artificial intelligence applications of machine learning and deep learning are promoting disaggregated architectures that allow system optimization by optical interconnection of cloud resources. Additional advantages of resource performance optimization and cost reduction by modularization and commoditization are expected. How do these trends affect the PD in the data center?

A finer grained version of system disaggregation is transceiver disaggregation: a remote laser provides stable, low relative intensity noise, high reliability performance. Modulators and detectors merge into processors to reduce metal interconnect length and associated parasitics. Monolithic integration of the PD in the upper levels of the processor interconnection stack is expected. The dimensions of optical devices are compatible with feature sizes and wire diameters at the upper levels, and the interconnection length is minimal (see Fig. 4.14). A second driver for BEOL integration is the retrofit of Ge absorbers into standard silicon circuit technology. In this case, both monolithic and wafer bonding technologies will be employed.

The data center system requirement for PD speed is the system line rate (component baud rate). A baud rate limit of ~ 100 Gb/s is determined today by constraints of metal interconnects at the chip and package levels. As monolithic integration of detectors with processors matures, higher baud rates are likely to enable package I/O scaling to 1 Pb/s with concomitant data center capacity scaling.

References

[1] L.C. Kimerling, J. Michel, Monolithic microphotonic integration on the silicon platform, ECS Trans. 41 (3) (2011) 3–13. Available from: https://doi.org/10.1149/1.3633281.

[2] J. Liu, R. Camacho-Aguilera, J.T. Bessette, X. Sun, X. Wang, Y. Cai, et al., Ge-on-Si optoelectronics, Thin Solid Films 520 (8) (2012) 3354–3360. Available from: https://doi.org/10.1016/j.tsf.2011.10.121.
[3] J. Liu, Monolithically integrated Ge-on-Si active photonics, Photonics 1 (3) (2014) 162–197. Available from: https://doi.org/10.3390/photonics1030162.
[4] H. Li, et al., Silicon waveguide integrated with germanium photodetector for a photonic-integrated FBG interrogator, Nanomaterials 10 (9) (2020) 1683. Available from: https://doi.org/10.3390/nano10091683.
[5] D. Ahn, C. Hong, J. Liu, et al., High performance, waveguide integrated Ge photodetectors, Opt. Express 15 (7) (2007) 3916–3921. Available from: https://doi.org/10.1364/OE.15.003916.
[6] L. Vivien, A. Polzer, D. Marris-Morini, et al., Zero-bias 40Gbit/s germanium waveguide photodetector on silicon, Opt. Express 20 (2) (2012) 1096–1101. Available from: https://doi.org/10.1364/OE.20.001096.
[7] A. Novack, M. Gould, Y. Yang, et al., Germanium photodetector with 60 GHz bandwidth using inductive gain peaking, Opt. Express 21 (23) (2013) 28387–28393. Available from: https://doi.org/10.1364/OE.21.028387.
[8] S. Liao, N.-N. Feng, D. Feng, et al., 36 GHz submicron silicon waveguide germanium photodetector, Opt. Express 19 (11) (2011) 10967–10972. Available from: https://doi.org/10.1364/OE.19.010967.
[9] D. Feng, W. Qian, H. Liang, B.J. Luff, M. Asghari, High-speed receiver technology on the SOI platform, IEEE J. Sel. Top. Quantum Electron. 19 (2) (2013) 3800108. Available from: https://doi.org/10.1109/JSTQE.2012.2213804.
[10] J. Michel, R. Wen, D. Ma, G. Roscioli, X. Zhao, Ge-based photonic devices for CMOS integration (Conference Presentation), in: G.T. Reed, A.P. Knights (Eds.), Silicon Photonics XIII, vol. 10537, SPIE, 2018. Available from: http://doi.org/10.1117/12.2289284.
[11] Y. Lin, D. Ma, K.H. Lee, et al., PIC-integrable, uniformly tensile-strained Ge-on-insulator photodiodes enabled by recessed SiN_x stressor, Photon Res. 9 (7) (2021) 1255–1263. Available from: https://doi.org/10.1364/PRJ.419776.
[12] L. Vivien, M. Rouvière, J.-M. Fédéli, et al., High speed and high responsivity germanium photodetector integrated in a silicon-on-insulator microwaveguide, Opt. Express 15 (15) (2007) 9843–9848. Available from: https://doi.org/10.1364/OE.15.009843.
[13] S. Assefa, F. Xia, Y.A. Vlasov, Reinventing germanium avalanche photodetector for nanophotonic on-chip optical interconnects, Nature 464 (7285) (2010) 80–84. Available from: https://doi.org/10.1038/nature08813.
[14] S. Assefa, F. Xia, S.W. Bedell, et al., CMOS-integrated high-speed MSM germanium waveguide photodetector, Opt. Express 18 (5) (2010) 4986–4999. Available from: https://doi.org/10.1364/OE.18.004986.
[15] K.S. Giboney, J.W. Rodwell, J.E. Bowers, Traveling-wave photodetector theory, IEEE Trans. Microw. Theory Tech. 45 (8) (1997) 1310–1319. Available from: https://doi.org/10.1109/22.618429.
[16] K. Sun, A. Beling, High-speed photodetectors for microwave photonics, Appl. Sci. 9 (4) (2019). Available from: https://doi.org/10.3390/app9040623.
[17] M. Piels, J.E. Bowers, 40 GHz Si/Ge uni-traveling carrier waveguide photodiode, J. Light. Technol. 32 (20) (2014) 3502–3508. Available from: https://doi.org/10.1109/JLT.2014.2310780.
[18] L. Bogaert, K. Van Gasse, T. Spuesens, G. Torfs, J. Bauwelinck, G. Roelkens, Silicon photonics traveling wave photodiode with integrated star coupler for high-linearity

mm-wave applications, Opt. Express 26 (26) (2018) 34763–34775. Available from: https://doi.org/10.1364/OE.26.034763.
[19] M. Huang, P. Cai, S. Li, et al., 56GHz waveguide Ge/Si avalanche photodiode, in: 2018 Optical Fiber Communications Conference and Exposition (OFC), 2018, pp. 1–3.
[20] J. Michel, S.J. Koester, J. Liu, X. Wang, M.W. Geis, S.J. Spector, et al., Photodetectors, Handbook of Silicon Photonics, CRC Press, 2013, pp. 479–552. Available from: http://doi.org/10.1201/b14668.
[21] N.M. Fahrenkopf, C. McDonough, G.L. Leake, Z. Su, E. Timurdogan, D.D. Coolbaugh, The AIM photonics MPW: a highly accessible cutting edge technology for rapid prototyping of photonic integrated circuits, IEEE J. Sel. Top. Quantum Electron 25 (5) (2019) 1–6. Available from: https://doi.org/10.1109/JSTQE.2019.2935698.
[22] A. Beling, J.C. Campbell, InP-based high-speed photodetectors, J. Light. Technol. 27 (3) (2009) 343–355. Available from: https://doi.org/10.1109/JLT.2008.2008399.
[23] Y. Kang, Z. Huang, Y. Saado, et al., High performance Ge/Si avalanche photodiodes development in intel, in: Optical Fiber Communication Conference/National Fiber Optic Engineers Conference 2011. Optical Society of America, 2011, OWZ1. Available from: https://doi.org/10.1364/OFC.2011.OWZ1.
[24] Y. Ishikawa, K. Wada, D.D. Cannon, J. Liu, H.-C. Luan, L.C. Kimerling, Strain-induced band gap shrinkage in Ge grown on Si substrate, Appl. Phys. Lett. 82 (13) (2003) 2044–2046. Available from: https://doi.org/10.1063/1.1564868.
[25] J. Liu, D.D. Cannon, K. Wada, et al., Deformation potential constants of biaxially tensile stressed Ge epitaxial films on Si(100), Phys Rev. B 70 (15) (2004) 155309. Available from: https://doi.org/10.1103/PhysRevB.70.155309.
[26] J. Liu, J. Michel, W. Giziewicz, et al., High-performance, tensile-strained Ge p-i-n photodetectors on a Si platform, Appl. Phys. Lett. 87 (10) (2005) 103501. Available from: https://doi.org/10.1063/1.2037200.
[27] J. Wei, F. Xia, S.R. Forrest, A high-responsivity high-bandwidth asymmetric twin-waveguide coupled InGaAs-InP-InAlAs avalanche photodiode, IEEE Photonics Technol. Lett. 14 (11) (2002) 1590–1592. Available from: https://doi.org/10.1109/LPT.2002.803894.
[28] S. Xie, S. Zhang, C.H. Tan, InGaAs/InAlAs avalanche photodiode with low dark current for high-speed operation, IEEE Photonics Technol. Lett. 27 (16) (2015) 1745–1748. Available from: https://doi.org/10.1109/LPT.2015.2439153.
[29] J. Michel, J. Liu, L.C. Kimerling, High-performance Ge-on-Si photodetectors, Nat. Photonics. 4 (8) (2010) 527–534. Available from: https://doi.org/10.1038/nphoton.2010.157.
[30] H.-C. Luan, D.R. Lim, K.K. Lee, et al., High-quality Ge epilayers on Si with low threading-dislocation densities, Appl. Phys. Lett. 75 (19) (1999) 2909–2911. Available from: https://doi.org/10.1063/1.125187.
[31] S.B. Samavedam, M.T. Currie, T.A. Langdo, E.A. Fitzgerald, High-quality germanium photodiodes integrated on silicon substrates using optimized relaxed graded buffers, Appl. Phys. Lett. 73 (15) (1998) 2125–2127. Available from: https://doi.org/10.1063/1.122399.
[32] Y. Ishikawa, K. Noguchi, M. Tachibana, et al., Selective epitaxy of submicron Ge wire structures for photodetectors and optical modulators in Si photonics, ECS Trans. 104 (4) (2021) 147–155. Available from: https://doi.org/10.1149/10404.0147ecst.
[33] S. Sonoi, R. Katamawari, M. Shimokawa, et al., Direct bandgap control by narrowing the germanium strip structure on silicon for C + L band photonic devices, IEEE J.

Quantum Electron. 58 (5) (2022) 8400209. Available from: https://doi.org/10.1109/JQE.2022.3203128.

[34] L. Vivien, J. Osmond, J.-M. Fédéli, et al., 42 GHz p.i.n germanium photodetector integrated in a silicon-on-insulator waveguide, Opt. Express 17 (8) (2009) 6252−6257. Available from: https://doi.org/10.1364/OE.17.006252.

[35] K.-W. Ang, T.-Y. Liow, M.-B. Yu, et al., Low thermal budget monolithic integration of evanescent-coupled Ge-on-SOI photodetector on Si CMOS platform, IEEE J. Sel. Top. Quantum Electron. 16 (1) (2010) 106−113. Available from: https://doi.org/10.1109/JSTQE.2009.2025142.

[36] N.-N. Feng, P. Dong, D. Zheng, et al., Vertical p-i-n germanium photodetector with high external responsivity integrated with large core Si waveguides, Opt. Express 18 (1) (2010) 96−101. Available from: https://doi.org/10.1364/OE.18.000096.

[37] S. Park, T. Tsuchizawa, T. Watanabe, et al., Monolithic integration and synchronous operation of germanium photodetectors and silicon variable optical attenuators, Opt. Express 18 (8) (2010) 8412−8421. Available from: https://doi.org/10.1364/OE.18.008412.

[38] H. Nishi, T. Tsuchizawa, R. Kou, et al., Monolithic integration of a silica AWG and Ge photodiodes on Si photonic platform for one-chip WDM receiver, Opt. Express 20 (8) (2012) 9312−9321. Available from: https://doi.org/10.1364/OE.20.009312.

[39] T. Hiraki, H. Nishi, T. Tsuchizawa, et al., Si-Ge-silica monolithic integration platform and its application to a 22-Gb/s x 16-ch WDM receiver, IEEE Photonics J. 5 (4) (2013) 4500407. Available from: https://doi.org/10.1109/JPHOT.2013.2269676.

[40] T. Yin, R. Cohen, M.M. Morse, et al., 31GHz Ge n-i-p waveguide photodetectors on silicon-on-insulator substrate, Opt. Express 15 (21) (2007) 13965−13971. Available from: https://doi.org/10.1364/OE.15.013965.

[41] G. Masini, S. Sahni, G. Capellini, J. Witzens, C. Gunn, High-speed near infrared optical receivers based on Ge waveguide photodetectors integrated in a CMOS process, Adv. Opt. Technol. 2008 (2008) 196572. Available from: https://doi.org/10.1155/2008/196572.

[42] Y. Ishikawa, S. Saito, Ge-on-Si photonic devices for photonic-electronic integration on a Si platform, IEICE Electron. Expr. 11 (24) (2014). Available from: https://doi.org/10.1587/elex.11.20142008.

[43] Y. Takada, J. Osaka, Y. Ishikawa, K. Wada, Effect of mesa shape on threading dislocation density in Ge epitaxial layers on Si after post-growth annealing, Jpn. J. Appl. Phys. 49 (4) (2010) 04DG23. Available from: https://doi.org/10.1143/jjap.49.04dg23.

[44] S. Nagatomo, Y. Ishikawa, S. Hoshino, Near-infrared laser annealing of Ge layers epitaxially grown on Si for high-performance photonic devices, J. Vac. Sci. Technol. B 35 (5) (2017) 51206. Available from: https://doi.org/10.1116/1.4995321.

[45] N. Higashitarumizu, Y. Ishikawa, Enhanced direct-gap light emission from Si-capped n + -Ge epitaxial layers on Si after post-growth rapid cyclic annealing: impact of non-radiative interface recombination toward Ge/Si double heterostructure lasers, Opt. Express 25 (18) (2017) 21286−21300. Available from: https://doi.org/10.1364/oe.25.021286.

[46] J.E. Ayers, The measurement of threading dislocation densities in semiconductor crystals by X-ray diffraction, J. Cryst. Growth 135 (1) (1994) 71−77. Available from: https://doi.org/10.1016/0022-0248(94)90727-7.

[47] G. Wang, R. Loo, E. Simoen, et al., A model of threading dislocation density in strain-relaxed Ge and GaAs epitaxial films on Si (100), Appl. Phys. Lett. 94 (10) (2009) 102115. Available from: https://doi.org/10.1063/1.3097245.

[48] M. Yako, Y. Ishikawa, K. Wada, Coalescence induced dislocation reduction in selectively grown lattice-mismatched heteroepitaxy: theoretical prediction and experimental verification, J. Appl. Phys. 123 (18) (2018) 185304. Available from: https://doi.org/10.1063/1.5011421.
[49] L. Giovane, H.C. Luan, A.M. Agarwal, L.C. Kimerling, Correlation between leakage current density and threading dislocation density in SiGe p-i-n diodes grown on relaxed graded buffer layers, Appl. Phys. Lett. 78 (2001) 541–543. Available from: https://doi.org/10.1063/1.1341230.
[50] C. Masini, L. Calace, G. Assanto, H.-C. Luan, L.C. Kimerling, High-performance p-i-n Ge on Si photodetectors for the near infrared: from model to demonstration, IEEE Trans. Electron. Dev. 48 (6) (2001) 1092–1096. Available from: https://doi.org/10.1109/16.925232.
[51] L. Colace, G. Assanto, Germanium on silicon for near-infrared light sensing, IEEE Photonics J. 1 (2) (2009) 69–79. Available from: https://doi.org/10.1109/JPHOT.2009.2025516.
[52] L. Colace, P. Ferrara, G. Assanto, D. Fulgoni, L. Nash, Low dark-current germanium-on-silicon near-infrared detectors, IEEE Photonics Technol. Lett. 19 (22) (2007) 1813–1815. Available from: https://doi.org/10.1109/LPT.2007.907578.
[53] B. Son, Y. Lin, K.H. Lee, Q. Chen, C.S. Tan, Dark current analysis of germanium-on-insulator vertical p-i-n photodetectors with varying threading dislocation density, J. Appl. Phys. 127 (20) (2020) 203105. Available from: https://doi.org/10.1063/5.0005112.
[54] D.D. Cannon, J. Liu, Y. Ishikawa, et al., Tensile strained epitaxial Ge films on Si(100) substrates with potential application in L-band telecommunications, Appl. Phys. Lett. 84 (6) (2004) 906–908. Available from: https://doi.org/10.1063/1.1645677.
[55] Y. Ishikawa, K. Wada, J. Liu, et al., Strain-induced enhancement of near-infrared absorption in Ge epitaxial layers grown on Si substrate, J. Appl. Phys. 98 (1) (2005) 13501. Available from: https://doi.org/10.1063/1.1943507.
[56] Y. Ishikawa, K. Wada, Germanium for silicon photonics, Thin Solid Films 518 (6, Suppl. 1) (2010) S83–S87. Available from: https://doi.org/10.1016/j.tsf.2009.10.062.
[57] K. Noguchi, M. Nishimura, Y. Tsusaka, J. Matsui, Y. Ishikawa, Enhancement of L-band optical absorption in strained epitaxial Ge on Si-on-quartz wafer: Toward extended Ge photodetectors, J. Appl. Phys. 128 (13) (2020) 133107. Available from: https://doi.org/10.1063/5.0023703.
[58] C.G. de Walle, Band lineups and deformation potentials in the model-solid theory, Phys. Rev. B 39 (3) (1989) 1871–1883. Available from: https://doi.org/10.1103/PhysRevB.39.1871.
[59] K. Takeda, T. Hiraki, T. Tsuchizawa, et al., Contributions of Franz–Keldysh and avalanche effects to responsivity of a germanium waveguide photodiode in the L-band, IEEE J. Sel. Top. Quantum Electron. 20 (4) (2014) 64–70. Available from: https://doi.org/10.1109/JSTQE.2013.2295182.
[60] J. Fujikata, M. Noguchi, K. Kawashita, et al., High-speed Ge/Si electro-absorption optical modulator in C-band operation wavelengths, Opt. Express 28 (22) (2020) 33123–33134. Available from: https://doi.org/10.1364/OE.405447.
[61] L. Barrutia, E. Barrigón, I. García, I. Rey-Stolle, C. Algora, Effect of Ge autodoping during III-V MOVPE growth on Ge substrates, J. Cryst. Growth 475 (10) (2017) 378–383. Available from: https://doi.org/10.1016/j.jcrysgro.2017.06.022.

[62] Jf Liu, K.Y. Ahn, C.Y. Hong, et al., Waveguide integrated Ge p-i-n photodetectors on a silicon-on-insulator platform, in: 2006 Optics Valley of China International Symposium on Optoelectronics, 2006, pp. 1−4. Available from: https://doi.org/10.1109/OVCISO.2006.302697.

[63] M. Beals, J. Michel, J.F. Liu, et al., Process flow innovations for photonic device integration in CMOS, in: J.A. Kubby, G.T. Reed (Eds.), Silicon Photonics III, vol. 6898, SPIE, 2008, pp. 31−44. Available from: http://doi.org/10.1117/12.774576.

[64] J. Liu, M. Beals, A. Pomerene, et al., Waveguide-integrated, ultralow-energy GeSi electro-absorption modulators, Nat. Photonics. 2 (7) (2008) 433−437. Available from: https://doi.org/10.1038/nphoton.2008.99.

[65] H. Chen, P. Verheyen, P. De Heyn, et al., Dark current analysis in high-speed germanium p-i-n waveguide photodetectors, J. Appl. Phys. 119 (21) (2016) 213105. Available from: https://doi.org/10.1063/1.4953147.

[66] P. Kuo, J.L. Hoyt, J.F. Gibbons, J.E. Turner, D. Lefforge, Boron diffusion in Si and Si 1 − x Ge x, MRS Proceedings, vol. 379, Cambridge University Press, 1995, p. 373. Available from: http://doi.org/10.1557/PROC-379-373.

[67] S. Uppal, A.F.W. Willoughby, J.M. Bonar, et al., Diffusion of boron in germanium at 800−900°C, J. Appl. Phys. 96 (3) (2004) 1376−1380. Available from: https://doi.org/10.1063/1.1766090.

[68] A. Dimoulas, P. Tsipas, A. Sotiropoulos, E.K. Evangelou, Fermi-level pinning and charge neutrality level in germanium, Appl. Phys. Lett. 89 (25) (2006) 252110. Available from: https://doi.org/10.1063/1.2410241.

[69] B. Yang, J.J. Lin, S. Gupta, et al., Low-contact-resistivity nickel germanide contacts on n + Ge with phosphorus/antimony co-doping and schottky barrier height lowering, 2012 Int. Silicon-Germanium Technol. Device Meet. (ISTDM) (2012) 1−2. Available from: https://doi.org/10.1109/ISTDM.2012.6222472.

[70] D. Feng, S. Liao, P. Dong, et al., High-speed Ge photodetector monolithically integrated with large cross-section silicon-on-insulator waveguide, Appl. Phys. Lett. 95 (26) (2009) 261105. Available from: https://doi.org/10.1063/1.3279129.

[71] C.-H. Park, M. Yako, K. Wada, Y. Ishikawa, D. Ahn, High concentration phosphorus doping in Ge for CMOS-integrated laser applications, Solid State Electron. 154 (2019) 43−49. Available from: https://doi.org/10.1016/J.SSE.2019.02.007.

[72] J. Liu, X. Sun, D. Pan, et al., Tensile-strained, n-type Ge as a gain medium for monolithic laser integration on Si, Opt. Express 15 (18) (2007) 11272−11277. Available from: https://doi.org/10.1364/OE.15.011272.

[73] X. Sun, Ge-on-Si Light-Emitting Materials and Devices for Silicon Photonics (PhD thesis). Available from: https://dspace.mit.edu/handle/1721.1/53254.

[74] J. Baek, B. Ki, D. Kim, et al., Phosphorus implantation into in situ doped Ge-on-Si for high light-emitting efficiency, Opt. Mater. Express. 6 (9) (2016) 2939. Available from: https://doi.org/10.1364/ome.6.002939.

[75] H. Pan, R. Takahashi, K. Takinai, K. Wada, Photoluminescence quenching effect by Si cap in n + Ge on Si, in: G.T. Reed, M.R. Watts (Eds.), Silicon Photonics X, vol. 9367, 2015. 93671E-1-7. Available from: https://doi.org/10.1117/12.2078593.

[76] N.A. DiLello, J.L. Hoyt, Impact of post-metallization annealing on Ge-on-Si photodiodes passivated with silicon dioxide, Appl. Phys. Lett. 99 (3) (2011) 33508. Available from: https://doi.org/10.1063/1.3615800.

[77] N.A. DiLello, D.K. Johnstone, J.L. Hoyt, Characterization of dark current in Ge-on-Si photodiodes, J. Appl. Phys. 112 (5) (2012) 54506. Available from: https://doi.org/10.1063/1.4749259.

[78] B. Son, Y. Lin, K.H. Lee, Y. Wang, S. Wu, C.S. Tan, High speed and ultra-low dark current Ge vertical p-i-n photodetectors on an oxygen-annealed Ge-on-insulator platform with GeOx surface passivation, Opt. Express 28 (16) (2020) 23978–23990. Available from: https://doi.org/10.1364/OE.398199.
[79] C. Gunn, CMOS photonics for high-speed interconnects, IEEE Micro. 26 (2) (2006) 58–66. Available from: https://doi.org/10.1109/MM.2006.32.
[80] Soitec Inc. Photonics-SOI-Soitec. Available from: https://www.soitec.com/en/products/photonics-soi, 2021.
[81] W. Zhang, M. Ebert, J.D. Reynolds, B. Chen, X. Yan, H. Du, et al., Buried 3D spot-size converters for silicon photonics, Optica 8 (8) (2021) 1102–1108. Available from: https://doi.org/10.1364/OPTICA.431064.
[82] S.H. Jeong, Y. Tanaka, K. Morito, Silicon photonics optical demultiplexer technology for WDM optical interconnects, in: 2015 Opto-Electronics and Communications Conference (OECC), 2015, pp. 1–3. Available from: https://doi.org/10.1109/OECC.2015.7340286.
[83] H.-C. Ji, K.H. Ha, I.S. Joe, S.G. Kim, K.W. Na, D.J. Shin, et al., Optical interface platform for DRAM integration, in: Optical Fiber Communication Conference/National Fiber Optic Engineers Conference 2011, OThV4, 2011, pp. 1–3.
[84] H. Byun, J. Bok, K. Cho, et al., Bulk-Si photonics technology for DRAM interface, Photon Res. 2 (3) (2014) A25–A33. Available from: https://doi.org/10.1364/PRJ.2.000A25.
[85] V. Stojanović, R.J. Ram, M. Popović, et al., Monolithic silicon-photonic platforms in state-of-the-art CMOS SOI processes, Opt. Express 26 (10) (2018) 13106–13121. Available from: https://doi.org/10.1364/OE.26.013106.
[86] L. Alloatti, S.A. Srinivasan, J.S. Orcutt, R.J. Ram, Waveguide-coupled detector in zero-change complementary metal–oxide–semiconductor, Appl. Phys. Lett. 107 (4) (2015) 41104. Available from: https://doi.org/10.1063/1.4927393.
[87] A.H. Atabaki, S. Moazeni, F. Pavanello, et al., Integrating photonics with silicon nanoelectronics for the next generation of systems on a chip, Nature 556 (7701) (2018) 349–354. Available from: https://doi.org/10.1038/s41586-018-0028-z.
[88] Y.H.D. Lee, M. Lipson, Back-end deposited silicon photonics for monolithic integration on CMOS, IEEE J. Sel. Top. Quantum Electron. 19 (2) (2013) 8200207. Available from: https://doi.org/10.1109/JSTQE.2012.2209865.
[89] K. Preston, S. Manipatruni, C.B. Poitras, M. Lipson, 2.5 Gbps electro-optic modulator in deposited silicon, in: 2009 Conference on Lasers and Electro-Optics and 2009 Conference on Quantum Electronics and Laser Science Conference, 2009, pp. 1–2. Available from: https://doi.org/10.1364/CLEO.2009.CTuBB1.
[90] K.A. McComber, X. Duan, J. Liu, J. Michel, L.C. Kimerling, Single-crystal germanium growth on amorphous silicon, Adv. Funct Mater. 22 (5) (2012) 1049–1057. Available from: https://doi.org/10.1002/adfm.201102015.
[91] B.G. Lee, A. Biberman, J. Chan, K. Bergman, High-performance modulators and switches for silicon photonic networks-on-chip, IEEE J. Sel. Top. Quantum Electron. 16 (1) (2010) 6–22. Available from: https://doi.org/10.1109/JSTQE.2009.2028437.
[92] A.V. Krishnamoorthy, Hybrid silicon photonics integration, in: Optical Fiber Communication Conference. OSA Technical Digest (online). Optical Society of America, 2014, Th4G.4. Available from: https://doi.org/10.1364/OFC.2014.Th4G.4.
[93] H.D. Thacker, X. Zheng, J. Lexau, et al., An all-solid-state, WDM silicon photonic digital link for chip-to-chip communications, Opt. Express 23 (10) (2015) 12808–12822. Available from: https://doi.org/10.1364/OE.23.012808.

[94] J.T. Bessette, D. Ahn, Vertically stacked microring waveguides for coupling between multiple photonic planes, Opt. Express 21 (11) (2013) 13580−13591. Available from: https://doi.org/10.1364/OE.21.013580.
[95] N. Nishiyama, J. Kang, Y. Kuno, et al., Si-photonics-based layer-to-layer coupler toward 3D optical interconnection, IEICE Trans. Electron. E101.C (7) (2018) 501−508. Available from: https://doi.org/10.1587/transele.E101.C.501.
[96] J. Kunstler, K. Ganguly, Transformation of wireline access networks requires smarter solutions, in: Ovum TMT Intelligence Informa, 2016.
[97] F. Villa, F. Severini, F. Madonini, F. Zappa, SPADs and SiPMs arrays for long-range high-speed light detection and ranging (LiDAR), Sensors 21 (11) (2021). Available from: https://doi.org/10.3390/s21113839.
[98] J.C. Campbell, S. Demiguel, F. Ma, et al., Recent advances in avalanche photodiodes, IEEE J. Sel. Top. Quantum Electron. 10 (4) (2004) 777−787. Available from: https://doi.org/10.1109/JSTQE.2004.833971.
[99] IEEE Standard for Ethernet—Amendment 10: Media Access Control Parameters, Physical Layers, and Management Parameters for 200 Gb/s and 400 Gb/s Operation. IEEE Std 8023bs-2017 (Amendment to IEEE 8023-2015 as Amend by IEEE's 8023bw-2015, 8023by-2016, 8023bq-2016, 8023bp-2016, 8023br-2016, 8023bn-2016, 8023bz-2016, 8023bu-2016, 8023bv-2017, IEEE 8023-2015/Cor1−2017). Published online 2017, pp. 1−372. Available from: https://doi.org/10.1109/IEEESTD.2017.8207825.
[100] Y. Kang, H.-D. Liu, M. Morse, et al., Monolithic germanium/silicon avalanche photodiodes with 340 GHz gain−bandwidth product, Nat. Photonics 3 (1) (2009) 59−63. Available from: https://doi.org/10.1038/nphoton.2008.247.
[101] D. Pan, New progress of Ge/Si APD/PD, in: CIOE Invited Talk, 2015.
[102] M. Huang, P. Cai, S. Li, et al., Breakthrough of 25Gb/s germanium on silicon avalanche photodiode, in: 2016 Optical Fiber Communications Conference and Exhibition (OFC), 2016, pp. 1−3.
[103] F. Chang, Y. Sun, R. Lingle, et al., First demonstration of PAM4 transmissions for record reach and high-capacity SWDM links over MMF using 40G/100G PAM4 IC chipset with real-time DSP, in: Optical Fiber Communication Conference. Optical Society of America, 2017, Tu2B.2. Available from: https://doi.org/10.1364/OFC.2017.Tu2B.2.
[104] K. Zhong, X. Zhou, Y. Wang, et al., Amplifier-less transmission of 56 Gbit/s PAM4 over 60 km using 25 Gbps EML and APD, in: Optical Fiber Communication Conference. Optical Society of America, 2017, Tu2D.1. Available from: https://doi.org/10.1364/OFC.2017.Tu2D.1.
[105] GR-468-CORE, Generic Reliability Assurance Requirements for Optoelectronic Devices Used in Telecommunications Equipment. Sept. 2004; issue 2. (Formerly Telcordia reliability assurance standards, https://telecom-info.njdepot.ericsson.net/site-cgi/ido/docs.cgi?ID = SEARCH&DOCUMENT = GR-468).
[106] B. Shi, F. Qi, P. Cai, X. Chen, Z. He, Y. Duan, et al., 106 Gb/s normal-incidence Ge/Si avalanche photodiode with high sensitivity, in: 2020 Optical Fiber Communications Conference and Exhibition (OFC). IEEE, Optica, 2020, M3D.2.
[107] C. Hong, B. Shi, F. Qi, P. Cai, Y. Duan, G. Hou, et al., High speed Ge/Si avalanche photodiode with high sensitivity for 50Gbit/s and 100Gbit/s optical access systems, in: 2022 Optical Fiber Communications Conference and Exhibition (OFC). IEEE, Optica, 2022, pp. 1−3.

[108] F. Nakajima, M. Nada, T. Yoshimatsu, High-sensitivity avalanche photodiode and receiver optical subassembly technology for 100-Gbit/s ethernet, NTT Tech. Rev. 13 (10) (2015) 1−7. Available from: https://www.ntt-review.jp/archive/ntttechnical.php?contents = ntr201510ra1.html.

[109] B. Chen, Y. Wan, Z. Xie, et al., Quantum dot avalanche photodetector on Si substrate, in: 2020 Conference on Lasers and Electro-Optics (CLEO), 2020, pp. 1−2.

[110] S. Luryi, Electronic devices using multilayer structures, in: P. Dhez, C. Weisbuch (Eds.), Physics, Fabrication, and Applications of Multilayered Structures, Springer, US, 1988, pp. 241−270. Available from: http://doi.org/10.1007/978-1-4757-0091-6_12.

[111] Y. Miyasaka, T. Hiraki, K. Okazaki, et al., Ge/graded-SiGe multiplication layers for low-voltage and low-noise Ge avalanche photodiodes on Si. Jpn. J. Appl. Phys. 55 (4S) (2016) 04EH10. Available from: https://doi.org/10.7567/jjap.55.04eh10.

This page intentionally left blank

Passive silicon photonic devices

5

Stefan Preble[1], *Barton Bergman*[2], *Lewis G. Carpenter*[2], *Lukas Chrostowski*[3], *Amit Dikshit*[2], *Michael Fanto*[4], *Wenhua Lin*[5], *Matthew van Niekerk*[1], *Mohammad Rakib Uddin*[2] and *Vijay Soorya Shunmuga Sundaram*[1]

[1]Rochester Institute of Technology, Rochester, NY, United States, [2]American Institute for Manufacturing Integrated Photonics (AIM Photonics), Albany, NY, United States, [3]The University of British Columbia, Vancouver, BC, Canada, [4]Air Force Research Laboratory, Rome, NY, United States, [5]Intel, Santa Clara, CA, United States

5.1 Introduction

Passive devices and circuits are the bedrock and framework of integrated photonic chips. They route, integrate, and interfere optical signals—forming the basis for all of the functionalities required for managing information with light, including: filters, multiplexers/demultiplexers, electro-optic modulators, photodetectors, and optical fiber couplers. The drivers for passive integrated devices in data communications are based around: cost to manufacture, component density, and waveguide propagation loss. Integrated photonic devices for datacom and long haul telecommunications share many similarities, but datacom devices demand lower device costs because of their shorter transmission distance and lifespan [1]. Silicon-on-insulator (SOI) by virtue of its ready availability in 150−300-mm wafer formats and its considerably smaller price tag [2−5] when compared to III−Vs like InP [6] make it an attractive platform for the passive routing of optical data. Further to this, component density, that is, the number of waveguides, active/passive components: lasers, modulators, photodetectors, etc., that can be fit into a single area without device performance degradation or crosstalk is directly related to the high refractive index contrast (Δn) between silicon and silicon dioxide. The higher the Δn, the tighter bend radii can be made and the more components can be squeezed into an area. Silicon photonics has already realized numerous demonstrations of thousands of photonic components on a single die [2], and an example is shown in Fig. 5.1. The high index contrast of SOI does come with a minor tradeoff as higher Δn waveguides do suffer from higher losses thus the photonic designer must balance component density and loss. However, through continual advancements in fabrication, low losses are being achieved in integrated silicon photonic platforms [7].

Silicon photonics is being driven by the explosive growth in modern data-intensive applications, which are poised to push data traffic into the so-called

Integrated Photonics for Data Communication Applications. DOI: https://doi.org/10.1016/B978-0-323-91224-2.00001-1
© 2023 Elsevier Inc. All rights reserved.

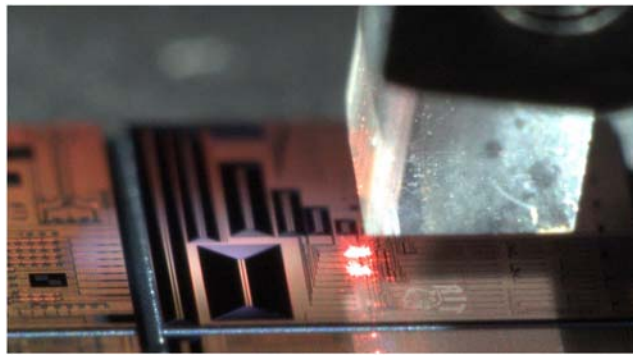

Figure 5.1 Photonic-integrated circuit chip (fabricated by AIM photonics and designed by the Rochester Institute of Technology) with passive device test circuits [waveguide spirals, directional couplers (DCs), Mach−Zehnder interferometers (MZIs), and more] being tested with a grating coupled optical fiber array.

"Zettabyte Era (ZB, 1 ZB = 0^{21} bytes)" [8]. These explosive data traffic demands will continue indefinitely over the foreseeable future and will force a paradigm shift in interconnects to keep pace [9]. Specifically, silicon photonic interconnects are replacing electrical interconnects in high performance computing and data centers [10]. To meet the interconnect bandwidth needs, there is continual drive to realize advances in circuit density, while realizing lower losses and higher reproducibility. This ends up putting stringent requirements on passive device performance.

In this chapter we will survey the key passive optical devices used in integrated photonic chips and compare the various approaches used to meet datacom application needs. We will begin with comparisons of: waveguides (Section 5.2.1), routing elements—bends/crossings/terminations (Section 5.2.2), splitters/combiners/couplers (Section 5.2.3), polarization diversity components (Section 5.2.4), and fiber-to-chip couplers (Section 5.2.5). In Section 5.3, we will focus on the primary application of passive circuits in datacom applications: wavelength filters and multiplexers/demultiplexers (WDM). We will summarize the common approaches for realizing wavelength division multiplexing and compare area, bandwidth, channel accuracy, insertion loss, and crosstalk. Lastly, at the end of this chapter, we will touch on future directions, including requirements, design/test methodologies, and the growing integration of other material systems with silicon photonics.

5.2 Survey of key passive optical devices

5.2.1 Waveguides

Waveguides operate using the optical principle of total internal reflection, which confines light into a high refractive index core, surrounded by a low refractive index cladding. The larger the difference between the core and cladding indices, the

smaller the waveguide can be, and the tighter the bends can be made. This is clearly seen in Table 5.1, which compares waveguides in a variety of platforms. There are two basic types of integrated photonic waveguides: strip and rib waveguides. Strip waveguides have cores that are fully etched in the manufacturing process, yielding the maximum confinement of light. In contrast, rib waveguides are only partially etched with the partially etched regions typically used to conduct electricity or heat when integrated into lasers, electro-optic modulators, photodetectors or phase shifters [11]. The lower optical confinement of rib waveguides can also be used to reduce loss.

The advantages of SOI are exemplified in Table 5.1. Both SOI strip and rib waveguides possess high Δn and tight bend radii enabling high levels of component density and relatively low propagation losses, generally <1 dB/cm [2,7,12,13] for state-of-the-art fabrication. A comparison in component density between conventional InP and SOI shows there is three orders of magnitude difference, ~ 1000 components/chip for InP [23] compared to ~ 1 million per square centimeter [24] for SOI due to the tighter bend radii of SOI waveguides. SOI is also readily integrated with SiN waveguides. SiN waveguides have achieved ultra-low losses [17] and can be evanescently coupled to silicon strip waveguides with ~ 0.1 dB insertion losses [4], enabling seamless transitions from silicon-to-SiN layers (Section 5.2.2.2).

5.2.1.1 Waveguides

Table 5.1 shows both silicon strip and rib waveguides. Typically, strip waveguides are used as the basis of passive device construction, see later Section 5.2.2. Owing to their rectangular shape, the mode shape is uniform and as a consequence has moderate interaction with the sidewalls. The high confinement of strip waveguides allow for tight bends, and low propagation loss can be achieved by using wide strip waveguides ($2-3$ μm). Alternatively, the rib waveguide can confine light more tightly within the core, achieving low losses because of the reduced interaction with the waveguide sidewall surface roughness. By tailoring the widths and heights of a rib waveguide, the photonic designer can engineer propagation loss, mode shape, and dispersion. The rib is used not only for chip routing but also the construction of active devices (see Chapters 2–4). For example, the ability for mode size engineering and placement has made the rib particularly useful for integrating areas of high doping near waveguiding regions while minimizing propagation losses for phase modulators based on both thermo-optic and plasma dispersion; namely the optical field is low in the thin etched region, therefore the electrical contacts and heavy dopants do not introduce significant loss.

The high density Si waveguides shown in Table 5.1 are based on SOI platforms on CMOS standard $200-300$ mm wafers. However, there are notable Si and SiN waveguide systems based on thick micron scale cores in Si [25] and SiN [15,16]. Waveguides in 3 μm thick silicon offer good propagation losses of $0.1-0.15$ dB/cm [25] in both Si rib and strip albeit for a 150 mm wafer size. Remarkably low losses of 0.01 dB have been achieved in 900 nm thick SiN [15,16] with notable dispersion properties allowing dissipative Kerr soliton generation in rings, but again these

Table 5.1 Comparison of integrated photonic waveguides at 1.55 μm.

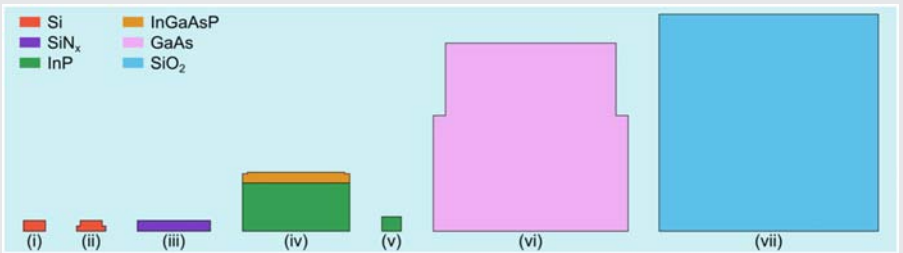

	SOI		**SiNx**	**InP**		**GaAs**	**SiO_2**
	(i) Strip	**(ii) Rib**	**(iii) Strip**	**(iv) Rib**	**(v) Strip**	**(vi) Rib**	**(vii) Strip**
Thickness (nm)	220−300	215−300	220	300−1500	300	1500−3900	4500
Width (nm)	300−500 (O- to C/L-band)	350−600	1500	2000	400	3500	4500
Effective index	2.4−2.7	2.5−2.8	1.5−1.8	3	-	2.5−3.5	-
$\Delta n = \frac{(n^2_{\text{ncore}} - n^2_{\text{clad}})}{2n^2_{\text{core}}}$	0.41	-	0.23	0.01	0.40	0.04	0.18
Bend radius (μm)	<10	>10	100	>1000	<10	>1000	>1000
Loss (dB/cm)	<2	<0.5	*0.08−0.5*	*0.45*	*2.5*	*1.6*	0.07
Wafer size (mm)	200−300	200−300	300	75	75	200	200
References	[2,7,12,13]	[2,13,14]	[7,15−17]	[18,19]	[20]	[21]	[22]

devices have only been demonstrated on small 100 mm wafer formats. Also, very thin SiN has been demonstrated with ultra low losses, for a core height of 65 nm and a width of 4.2 μm, propagation losses of 0.03 dB/cm [17] have been shown.

The SOI platform, combining both Si and SiN waveguides, allows for the transmission of both near infrared (NIR) and visible light. The cores have transmission windows from 1.2−8 and 0.4−4.6 μm [26] for Si and SiN, respectively. However, MIR (mid-infrared) transmission is curtailed to wavelengths <3.6 μm, where absorption significantly increases at longer wavelengths in the silicon dioxide cladding typically used in these platforms due to strong absorption caused by water in the material. Power handling in the SOI platform is quite different for Si versus SiN owing to the nonlinearities present in Si. For Si strip waveguides (145 nm × 540 nm), optical power of up to $\sim$200 mW was transmitted, but nonlinear absorption is observed at $\sim$150 mW for both 300 and 500 μm long waveguides [27]. In contrast, for thick SiN waveguides (820 nm × 1500 nm), $\sim$1 W level pumping has been achieved with sustained ring resonator Qs of 5×10^6 [28].

5.2.1.2 Manufacturability (loss and dimensional control)

Propagation losses in integrated waveguides are extensively studied but less so is the complex interplay with fabrication, integration, and the resultant device performance of both passives and actives. The photonic designer and process integrator both need to work in unison to reduce optical scatter and absorption while designing for manufacturing to ensure low loss operation and good performance. With regard to waveguide propagation loss and the reduction of optical scattering, waveguide surface roughness and particularly sidewall surface roughness need to be minimized. Often by careful selection of photolithography and etching parameters this can be achieved [12,13]. Typically, to achieve state-of-the-art propagation losses, 193 nm immersion lithography is used to produce both Si and SiN waveguides in 200/300-mm wafer formats with the photolithography yielding critical dimensions of 100 nm [2,4,12,13]. Another key device parameter is waveguiding film thickness, which is typically the starting SOI silicon and has thickness variation for 200 mm SOI wafers of 0.64 nm 1σ, with a 2.8 nm range for 220 nm thick films [2]. Etch control in terms of width, depth (for ribs) and sidewall angle and their variation are critical to propagation loss, and final etch depth variations in silicon of 2.5−3.3 nm has been demonstrated [5]. Annealing treatments utilizing hydrogen chemistries have proven effective at reducing losses in both strip and rib Si waveguide losses [13]. However, care must be taken as hydrogen causes absorption in the C-band for SiN [29] and at 1400 nm for SiO_2 cladding [30].

The reduction of material absorption loss is also critical for reducing propagation loss in integrated waveguides. The incredibly low losses achieved by standard telecommunications single mode optic fiber optic shows the importance of material purity in planar platforms, in particular for ring resonators and particularly SiN [17,29]. Like the germanosilicates used in fibers, SiN waveguide absorption losses are also sensitive to transition metals, such as Fe, Cu, and Cr [29]. Strict CMOS contamination control can be used to reduce the ingress of these metals in both Si

and SiN waveguides, by regularly monitoring physical and chemical contamination by spectroscopic/fluorescence measurements [4]. Furthermore, migration of active device dopants and metals into waveguiding regions can be minimized by considering thermal budgets and by the simple addition of diffusion barriers, such as TiW or TiN [31]. Further to this, the addition of active devices along with their precision doping profiles further complicates integration and explains why passive integrations often have lower optical losses [7,12,32].

5.2.2 Routing elements

Silicon photonics excels at the routing of optical signals across a chip. The large refractive index contrast of strip waveguides enables compact bends and crossings, and the effective index can be controlled over a wide range for terminating optical circuits. In this section, the common routing elements (Table 5.2) will be described.

5.2.2.1 Bends

PIC layouts need straight waveguides as well as waveguide bends or curves of varying radii to efficiently route optical signals or create components, such as Y-Splitters. In addition to the waveguide propagation loss, bends have additional sources of loss, such as radiation loss, due to finite curvature and mode-mismatch at the junction between straight and bent sections. Designing waveguide bends involves a trade-off between loss and minimum radius. The dominant source of loss is mode-mismatch loss for TE mode and radiation loss for TM mode [11]. Typical loss for a 500-nm width, 220-nm height, 2-μm buried oxide (BOX) cladding, silicon strip waveguide is $<$0.01 dB/90° for a 5-μm radius [33] for TE mode and increases exponentially with decreasing radius [34]. Although this value of loss might seem insignificant, a modern PIC can contain hundreds of bends in a single path and so the total loss will quickly add up.

Circular 90° bends (quarter part of a circle) are not the optimum design choice for minimum loss, and adiabatic bends based on Bezier curve [35] or Clothoid curve (Euler bends) [33,36] have shown to have significantly lower loss for a given (average) radius of curvature due to reduced mode-mismatch loss. Combining Bezier or Euler bends with widening width along the bend lowers radiation loss, further reducing the total loss [33]. For a given radius of curvature and footprint, advanced techniques that require special fabrication steps, such as inverse design [37] and use of multilayer metamaterials (“eskid” technique) [38], have shown promising results for sharp bends.

5.2.2.2 Waveguide crossings and escalators

Waveguide crossings are four-port symmetric devices used to route an optical signal across an intersection with another waveguide that is usually on the same layer. As the density of components in PICs increases, 100 s of waveguide crossings have been used in a single design for applications, such as optical switches and quantum

Table 5.2 Routing elements (bends, crossings/escalators and terminations).

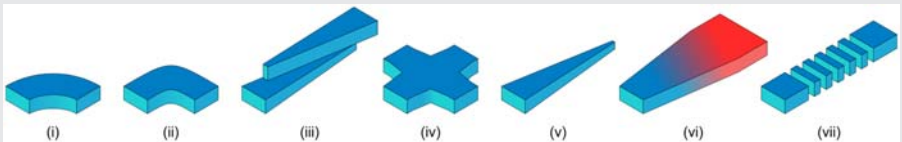

	Bends		**Crossings/escalators**		**Terminations**		
Types	Radial (i)	Adiabatic/ Bezier (ii)	Tapers (iii)	Multimode interference (MMI) or multimode waveguides (iv)	Tapers (v)	Absorbers (vi)	Reflectors (Bragg/ Sagnac) (vii)
Footprint	>10 μm (TM) <5 μm (TE)	>10 μm (TM) <3 μm (TE)	Large	Small	Large	Small−medium	Large
Cause of loss	Transition/ junction/ mode mismatch	Radiation	Scattering	Transition	N/A	N/A	Bragg—scattering Sagnac—splitter transitions
Cause of crosstalk	Reflection from transition	Scattering	Interfaces	Reflection and interfaces	Unguided light	None (with sufficient absorption)	Scattered light
Bandwidth	Large	Large	Large	Medium	Large	Large	Large

computation [39]. In addition, PICs that use edge-coupling [40] rely on low-loss and low crosstalk waveguide crossings for efficient layout. Although the insertion loss of a single waveguide crossing can be <0.1 dB [41], optical signals on a dense PIC can attenuate considerably after just a few crossings.

In addition to insertion loss, other parameters of interest for waveguide crossings are device footprint, crosstalk, back-reflection and sensitivity to random fabrication variations. Similar to most passive optical devices, waveguide crossing design involves trade-off between insertion loss and device footprint.

Crossings constructed from waveguides intersecting at 90° are not preferred since the abrupt intersection causes diffraction and excites higher order modes leading to degraded characteristics [41]. To avoid reflections and generation of higher order modes at the intersection, optimized tapers intersecting each other with slowly varying widths compared to the mode evolution are commonly used. For a typical optimized taper waveguide crossing fabricated in a standard CMOS compatible 300 mm foundry process with 2 μm BOX, 220 nm Si thickness and 480 nm width, the insertion loss for TE mode at 1550 nm is $\sim$0.2 dB after considering the impact of manufacturing variations as shown in Fig. 5.2 [42]. The footprint for this crossing is 8x8 μm^2 and the cross talk is -56 dB at 1550 nm. Further optimizations to

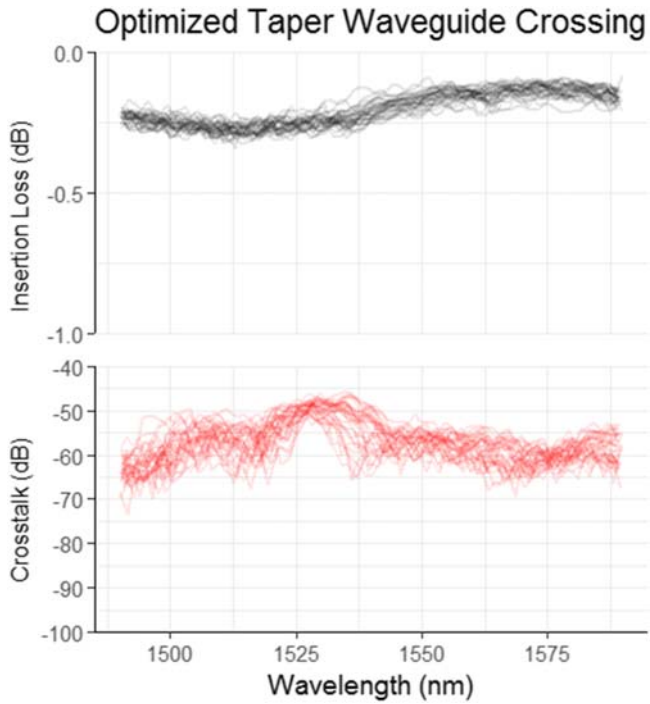

Figure 5.2 Optimized taper waveguide crossing insertion loss (top plot) and crosstalk (bottom plot).

this approach include double etched low-index contrast section [43] and optimized individual taper sections [44].

Another commonly used design technique for waveguide crossings is based on intersecting multimode interferometer (MMI) sections, where the length is chosen such that the self-image of the input mode is reproduced at the opposite port. Although MMIs are robust against fabrication variations, their footprint is comparatively higher. MMI-based crossings with a very low insertion loss of 43 mdB, device footprint of $14.3 \times 14.3\ \mu m^2$, and cross talk of -50 dB have been demonstrated using custom E-Beam lithography process [39]. The authors also further measure a backscatter of -44 dB in their design [39]. Performance of MMI-based crossings can be further improved by using rib waveguides [45], multisection tapers [46], and additional Optical Proximity Correction techniques [47], although these techniques are not readily compatible with reliable high-volume foundry manufacturing due to requirements of minimum feature size or additional postprocessing.

In addition to optimized taper and MMI-based waveguide crossings, designs using several advanced techniques have been demonstrated, which rely on precise fabrication control. Crossings using subwavelength gratings [48] that show a very low insertion loss of 23 mB have been demonstrated where the impact of loss due to diffraction is reduced. In addition, inverse design techniques have been applied for design optimization, and a significantly lower footprint of $4.5\ \mu m^2$ was achieved with an insertion loss of 0.2 dB [49]. Recently Maxwell fisheye lens has been used as a waveguide crossing due to its promise of low loss and ultrawide bandwidth covering multiple telecommunication bands [50]. For mode-division multiplexing applications, MMI-based multimode rib waveguide crossings with etch-depth optimization have been demonstrated with insertion loss of 0.1 dB (TE_0) and 0.2 dB (TE_1) [51].

Lastly, crossings can be achieved by making use of the third dimension, creating 2.5D/3D integrated photonic circuits that utilize multiple waveguiding layers [4,7,52]. Transferring light between layers is achieved with escalators [53–55], which adiabatically taper waveguides in each layer and utilize evanescent coupling to transfer light from layer to layer with low loss. As long as light is highly confined within each layer, crossings can be achieved with very low crosstalk.

5.2.2.3 Terminations

Photonic circuits need proper terminations to minimize unintended interference. There are two types of terminations: (1) reflectionless and (2) reflective. Reflectionless terminations are used at the end of circuits with no natural output signal. For example, unused ports of an interferometer or resonant circuit are common in multiplexing/demultiplexing systems. The unused ports must be properly terminated to avoid stray signals from reflecting back into the circuit, which can cause interference and crosstalk. The two common approaches for providing reflectionless terminations are the use of inverse tapers which adiabatically reduce the effective index to the surrounding cladding ($n_{eff} \sim n_{SiO2} = 1.45$) and/or absorbing the light using heavily doped silicon/germanium/metals. Absorbers are preferred

since simply reducing the effective index will result in unguided light that could cause crosstalk in other parts of the chip.

Reflective terminations are used when feedback is required, such as in resonant cavities. For example, they are used in external cavity lasers where a silicon photonic PIC provides feedback to a hybrid/heterogeneously integrated III−V waveguide amplifier [56,57]. The two most common approaches for reflecting light are Sagnac loops [58] or Bragg gratings [59]. Sagnac loops are simple to realize using a simple y-branch splitter/combiner by just connecting the two outputs of the branch together, forming a loop. The disadvantage of a Sagnac loop is the y-branch has a significant insertion loss, which will reduce the effective reflectivity. Bragg gratings are realized by periodically modulating the effective index of the waveguide (e.g., by widening/narrowing the waveguide). They can achieve high reflectivities and low loss but require high resolution photolithography and tight fabrication process control.

5.2.3 Splitters/combiners/couplers

Splitting and combining light are critical functions for building numerous optical circuits, such as interferometers, resonators, and optical switches. These operations are realized in passive integrated photonics with specific devices, namely y-branches (Section 5.2.3.1), MMIs (Section 5.2.3.2), DCs (Section 5.2.3.3), and adiabatic couplers (Section 5.2.3.4), as shown in Table 5.3. These devices share the common feature of input-to-output design versatility, ranging from standard input/output designs 1×2, 2×2, and 4×4 to more general $N \times M$ (where N and M are integers) designs.

Table 5.3 Splitters/combiners and couplers.

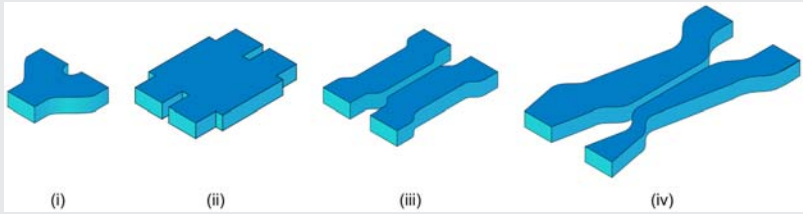

	Splitters/combiners		Couplers	
Types	(i) Branch	(ii) Multimode interference (MMI)	(iii) Directional	(iv) Adiabatic
Footprint	Small	Large	Medium	Large
Loss	High	Low	Low	Low
Bandwidth	Large	Medium	Low	High
Crosstalk/return loss	Large	Medium	Small	Small

5.2.3.1 Branch

The branch (or 1×2 splitter, y-branch) is useful for splitting or combining a signal with a small footprint, low-loss (<0.3 dB) and low splitting imbalance (<1 dB) [60,61]. Naive design, rigorous theoretical modeling, numerical optimization, and shape topology optimization are techniques used to create 1×2 splitters. The primary challenge with y-branches is the small footprint, where the transition in and out of the splitting/combining region is met with abrupt interfaces. These interfaces can cause reflection, insertion loss, and imbalance [61]. New design techniques, such as inverse design, are promising avenues toward improving the state-of-the-art performance, increasing resiliency to manufacturing variations and process differences [62].

5.2.3.2 Multimode interferometers

MMIs are multiport interference devices that allow for splitting/combining of the inputs/output signals through self-imaging. Structurally, the MMI has N excitation ports (generally single-mode inputs), a wide interference section, and M output ports. The wide interference section allows the excitation of higher-ordered modes from single-mode inputs, which enable restricted interference that creates a repeated "self-image" of the input at a characteristic length based on the material index of refraction and width. The excitation ports are commonly tapered to a slightly wider width, with tapers long enough to preserve single-mode excitation to decrease the device's insertion loss at the cost of device footprint. Arbitrary splitting ratios are achievable depending on the location and number of the input and output ports. Simple MMIs have a roughly broadband spectral response and are generally tolerant to minor manufacturing variations. These characteristics make them helpful routing elements for many circuit applications - particularly in high-index contrast platforms.

In addition, MMIs are applicable in coherent receiver circuits (Chapter 6), as the receiver or a component in the receiver circuit [63]. As a receiver, the MMI is generally designed with inputs/outputs of 2×4 or 4×4. The MMI mixes the local oscillator and the S/Rx signal for a coherent receiver, producing four different outputs, namely $I+$, $I-$, $Q+$, and $Q-$. These outputs are detected, forming classical I-Q values recovered from the original signal. Using the MMI, in Ref. [63], a low phase error (3 degrees) and a high common-mode rejection ratio of 30 dB were demonstrated

5.2.3.3 Directional

A traditional DC consists of two parallel waveguides, allowing controlled coupling of the excitation of the lowest ordered modes (symmetric and antisymmetric) in the coupled waveguide region. This coupling between waveguides is controlled through the distance between them (gap) and the interaction length. DCs are therefore simple to design for various conditions, such as 50/50 couplers or 1%/10% optical power taps for monitoring. However, the coupling behavior of DCs is generally

wavelength sensitive. For example, a 50/50 coupler will drift away from this designed ratio if we consider any wavelength deviation from the design. In addition, the coupling region is susceptible to manufacturing errors in the width of the individual waveguides or the gap of the coupling region. The wavelength sensitivity can be improved by dispersion engineering of the waveguides using subwavelength structures. The bandwidth was extended to 100 nm for a 50/50 splitter, albeit with higher insertion loss [64]. A typical application of a DC is implementation as the coupling mechanism for ring resonators, where a coupler with low loss and low coupling coefficient is often required.

5.2.3.4 Adiabatic

Adiabatic structures rely on a slow mode evolution to redistribute the power in the waveguides, leading to nearly lossless devices. These devices are insensitive to fabrication variations and offer broadband operation. However, they tend to be long (hundreds of microns) to ensure that the mode evolution is nearly lossless. A key challenge is reducing the device's footprint while maintaining low insertion loss and large operational bandwidth. Footprint reduction can be realized using subwavelength structures to increase the coupling between the waveguides. A 35-μm-long adiabatic splitter using subwavelength structures achieved an insertion loss of <0.11 dB over a 180-nm wavelength range [65]. Rapid adiabatic design optimizes the modal transition regions by speeding up or slowing down the geometric changes along the propagating direction. This method has been used to design couplers with very even power splitting ratios of $50\% \pm 1.08\%$ over 100 nm and insertion loss <0.2 dB in the O band [66,67].

5.2.4 Polarization diversity

Due to the birefringence of silicon, most components on the SOI platform are highly polarization sensitive. If the polarization of the incoming light is to be retained, different components are designed for optimal performance at each polarization and a polarization beam splitter (PBS) is employed to divert the incoming light toward appropriate circuits. Alternatively, a polarization rotator can be used to convert the incoming light to a single polarization, to be then manipulated on common components.

5.2.4.1 Polarization beam splitters

PBS are commonly designed based on two principles—Interferometry and Mode-Coupling. Interferometry based splitters are realized using MMIs or asymmetric MZIs. Conventional MMI-PBS are typically long (~ 1100 μm) since it requires the MMI length be a common multiple of the beat lengths of the TM and TE modes [68]. The length can be shortened using quasiimaging stage effect [69], using low-index guiding structures [70] or by cascading MMIs. The shortest MMI-based PBS was demonstrated using nanopatterning based on a nonlinear optimization algorithm

with a device footprint of just 2.4 μm × 2.4 μm [71]. Asymmetric MZI-PBS use the birefringence of the MZI arm waveguides for polarization splitting [72].

Mode-coupling-based devices rely on the different phase matching conditions of the fundamental transverse-electric (TE0) and transverse-magnetic (TM0) modes. This is done in one of two ways—(1) using a symmetric DC where the waveguide width is chosen such that the even and odd coupling lengths of the TE and TM modes are equal, shown in [73] or (2) using an asymmetric DC, which allows for only one polarization to strongly couple while the other remains unaffected. This is achieved using highly polarization-dependent systems, such as bent couplers [74] or systems with strong birefringence, such as single or multislot waveguides [75] or cross-slot waveguides [76]. Since they rely on phase matching, DC-based PBS are usually very sensitive to fabrication tolerances and are quite narrowband. Their typical lengths can vary on the orders of 10^1-10^4 μm.

5.2.4.2 Polarization rotators

Polarization rotators use phase-matching or mode-evolution to convert between TE and TM modes. In either case, the orthogonality of the TE and TM modes must be broken and this is done by breaking the symmetry of the effective index profile both horizontally and vertically. Horizontal symmetry is broken by varying the geometry of the waveguides gradually using tapers [77]. Adding etched layers or having an air cladding on top breaks the vertical symmetry.

DCs with lateral tapered sections with widths chosen for phase matching between TE and TM modes have been shown to rotate polarizations by direct TM0→TE0 rotation [78] or by TM0→TE1 evolution followed by TE1→TE0 phase matching [79]. The coupler lengths can be shortened using a dual-core Si_3N_4-Si waveguide for stronger birefringence [70].

In the case of mode evolution-based rotators, the polarization of the light changes as the geometry of the waveguide is gradually varied. Since the cross-section varies in the vertical direction as well, this method is harder to fabricate in the top-down process but the devices have a much smaller footprint compared to the DC-based rotators. Dual-layered waveguides [80] and waveguides with trenches [81] have been shown to produce ultra compact rotators. This can also be enhanced by tapering each layer [82] and by using layers of different materials [83]. A bi-level tapered section can also be used to rotate and act as a 50−50 splitter similar to a microwave magic-T with adiabatic performance [84]. Combinations of these methods have been employed to further shorten device footprints and minimize crosstalk using multimode waveguides and cascading adiabatic and bent couplers [85,86]. Rapid adiabatic couplers designed using parameterized functions to suppress unwanted mode conversions [87] have also been used in combination with the Magic-T and tapered waveguides to realize ultra-compact footprints (100 μm × 2 μm) with record low crosstalk < -14 dB [66].

5.2.5 Fiber-to-chip couplers

Optical fiber coupling is one of the most important and challenging aspects of integrated photonic chips. There is a large mismatch in the optical mode between standard single mode optical fibers (with a mode field diameter of approximately 10 μm) and high index contrast waveguides on silicon photonic chips, necessitating the need for novel mode converters that bridge between the fiber and chip scales. There are three primary approaches that have been developed, each with advantages and disadvantages (Table 5.4): (1) Edge coupling: couples light in/out of the edge of the chip, typically directly into fibers butted up into edge. The optical mode is expanded at the edge of the chip, which must be prepared to high optical quality using either polishing, deep plasma etching or cleaving. Low losses can be achieved with suitable mode converter design and chip fabrication processes. However, the density of optical connections is limited by the optical fiber spacing (typically 127−250 μm). (2) Grating coupling: couples vertically out of the surface of the chip using a diffraction grating. It requires the fibers to interface with the surface of the chip either directly vertical or with the use of mirrors (reflective or refractive). Losses are typically high but can be low with complicated chip fabrication processes that incorporate mirrors or control the interference. An advantage of grating coupling is that coupling can be realized anywhere on the surface of the chip, in 2D arrays enabling coupling to dense arrays of optical fibers or multiple fiber cores in multicore fibers. Grating coupling is also advantageous for wafer-level testing.

Table 5.4 Fiber-to-chip couplers.

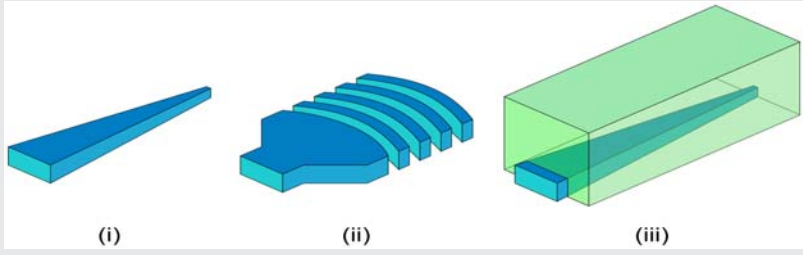

Types	Edge	Grating	Evanescent
Footprint	Medium to large	Small	Large
Loss	Low	High to low	Ultra-low
Alignment tolerance	Poor to good	Good	Poor vertically but forgiving in the same plane
Polarization-dependent loss	Good to excellent	Poor	Good
Bandwidth	Large	Low	Medium
Density	Low	Low−medium	Medium−high

(iii) Evanescent coupling: The optical mode is transferred from the chip into a polymer waveguide which can reside on another substrate or bridge from chip to fiber/chip to chip/chip to laser using directly laser written polymerization [photonic wire bonds (PWBs)]. It requires specialized polymer waveguide materials with suitably low losses and specialized processes and coupler designs that evanescently transition the optical mode from the polymer waveguide to the silicon (or silicon nitride) waveguides on the chip. Evanescent coupling can achieve very high coupling efficiencies provided the transitions are adiabatic and avoid scattering at interfaces. The use of polymer waveguides also enables a high density of connections that fan-in from an array of optical fibers to a dense spacing of waveguides on the chip.

5.2.5.1 Edge couplers

Edge couplers expand the waveguide optical mode at the edge of the chip to match the mode field diameter of optical fibers, which is on the order of 5−10 μm in diameter. Simply just widening a silicon waveguide to that width does not work because the effective index will increase to that of a 1D slab waveguide mode (approximately $n_{\text{eff}} \sim 3$ for 220-nm-thick silicon) resulting in a significant reflection. Even worse a 5−10-μm-wide waveguide has a highly elliptical mode with very little light above/below the waveguide, resulting in very poor modal overlap with circular gaussian optical fiber mode. To overcome these issues, the concept of the inverse taper [88,89] was invented and demonstrated. Instead of widening the waveguide at the edge of the chip, an inverse taper coupler is narrowed to a point with an effective index as closely matched to the optical fiber as possible ($n_{\text{eff}} \sim 1.5$). Consequently, the optical mode is also expanded laterally and vertically with modal overlaps can approach near-unity, enabling edge couplers approaching 1 dB coupling loss [90−92].

To further enhance the matching of the optical modes, additional low-index waveguides, based on either polymers or silicon-oxy-nitrides [90,93−95], are incorporated into the design. The mode matching is also highly dependent on the taper dimension, consequently, subwavelength metamaterial based designs have been employed [92,96] to realize couplers that are less immune to fabrication variation.

The high refractive index of the silicon substrate also has a significant impact on the mode matching. Specifically, as the mode expands it interacts with the substrate which distorts the mode, decreasing coupling efficiency. Consequently, substrate removal [92,97] or out-of-plane edge couplers [98] designs are being used to mitigate this. Lastly, the edge of the chip must be suitably prepared to ensure that the optical fiber can be butted against the chip. The large diameter of the fiber (125 μm) requires an edge that is highly vertical over at least 62.5 μm and is smooth. Previously, this required preparation of the edge of the chip using polishing—which is challenging to do in a manufacturing environment. Consequently, wafer scale definition of the edge is now common place where photolithography and deep plasma etching through the oxide and substrate ("Bosch Process") is used to create an optical quality edge for optical fiber coupling [92,99,100].

5.2.5.2 Grating couplers

Grating couplers use a periodic structure, on the order of the wavelength of light, to diffract light in/out of an optical fiber vertically placed over the grating coupler. It essentially directs light vertically incident on the chip into the waveguides using the principle of diffraction. A basic grating coupler is realized by etching a periodic pattern into the waveguide, either through the full depth of the waveguide or only partially [11,101–103]. The coupling efficiency, bandwidth of operation, angle of incidence (relative to vertical), and reflections are determined by the period/duty cycle of the pattern, waveguide thickness, etch depth, and surrounding dielectric thicknesses. However, it is very challenging to achieve coupling efficiencies of >50% with a basic grating coupler comprised of a single waveguide layer because at least half of the light will be diffracted down into the substrate.

To achieve higher coupling efficiencies, reflectors using either metal mirrors or periodic layers of dielectric (Bragg mirror) have been utilized to reflect the light going down into the substrate back into the optical fiber [104,105]. However, the integration of reflectors into the fabrication process is a significant manufacturing challenge. Another approach is to use multiple dielectric layers with varying refractive indices (i.e., SiN/Si/amorphous Si/Poly Si) and/or differing periodic patterns to destructively interfere the light that diffracts/reflects downwards into the substrate [106–109]. This approach has been very successful in directing a majority of the light vertically out-of-the-plane of the chip, achieving >75% coupling efficiencies to optical fibers. It is also compatible with PIC fabrication processes utilizing multiple waveguide layers or CMOS [4,106].

Another challenge to achieving high efficiencies is mode matching. The diffracted light needs to create a pattern that has high modal overlap with the gaussian optical fiber mode. This can be achieved using more complex apodized structures [110,111]. The size of the grating must also be as large as the optical fiber core, requiring long tapers down to the single mode waveguides or focusing designs that shorten the taper length [11,111]. It is also desirable to have grating couplers that operate with both polarizations of light. The high birefringence of integrated photonic waveguides makes it challenging to realize a single design that operates simultaneously for both polarizations. A variety of approaches have been taken to overcome this, including, two-dimensional grating patterns [112–114] or multilayer designs [115].

5.2.5.3 Evanescent couplers

To realize the highest coupling efficiencies evanescent coupling is employed where light is adiabatically transferred from one mode to another through the evanescent field. Specifically, light can be coupled from a waveguide on one substrate to another by placing the two waveguides in close contact together. With suitable mode design, typically by tapering the waveguide dimensions to match the effective indices of the two waveguides, it is possible to realize near unity coupling efficiency. The drawback of evanescent coupling is that it requires intimate contact

and/or tight control between the two waveguides to ensure that the evanescent modes overlap. It also requires ultra-long lengths to ensure adiabaticity, although new approaches for rapid adiabatic coupling are emerging at the chip scale [87].

Evanescent coupling is very forgiving to lateral misalignments, in turn, enabling passively aligned coupling. Relatively low loss (<2 dB), passively aligned fiber to PIC coupling has been demonstrated using a compliant polymer waveguide interface fiber optic connector where the optical mode is adiabatically transitioned from a polymer waveguide (butt-coupled to the optical fibers) into an inverse tapered waveguide on the top of the PIC chip [116,117]. A similar approach can be used for board-level interconnections between PIC chips [118]. Evanescent coupling directly from the fiber itself has also been demonstrated by tapering and capping an optical fiber with a high refractive index polymer, achieving 1-dB coupling loss [119]. Lastly, 3D-printed polymer waveguides are being used to realize PWB connections between fiber-to-PIC [120,121]. When evanescently coupled to a PIC waveguide the loss approaches less than 1 dB and since the PWB is 3D printed, the connections can be tailored to overcome minor misalignments between components.

5.3 Multiplexers/demultiplexers and filters

5.3.1 Requirements and current trends

Wavelength division multiplexing (WDM) is a key technology in optical fiber communication. It is commercially deployed to increase the capacity of fiber optic backbones, data center interconnects, and long distance data transmission. Optical multiplexer (Mux) and demultiplexer (Demux) are key components in WDM systems. To send multiple wavelength lanes down a single optical fiber, the wavelengths must be multiplexed (combined) by a Mux at the transmitting fiber end and de-multiplexed (separated) by a Demux at the receiving fiber end. Alternatively, a banded or a single wavelength add-drop filter selects one band of wavelengths or one wavelength and adds it into or drops it out of the optical lanes. WDM Mux/Demux component technology has been successfully developed, deployed, and proliferated in the coarse wavelength division multiplexing (CWDM) systems for intra/inter data centers (20 nm channel spacing, 4−8 channels), LAN WDM system for data center interconnects (4.5 nm or 800 GHz channel spacing, 4−8 channels), and noncoherent and coherent DWDM systems for metro/long hauls networks (50/100 GHz or 75 GHz dense channel spacing, 16−96 channels).

WDM systems invariably require good optical filters that meet stringent transmission system requirements. The data center systems must be low cost, low power, and high volume. For relatively short-reach intra data center applications, CWDM solutions are widely deployed, and the associated Mux/Demux components not only are high volume manufacturable, reliability, and low cost but also can have relaxed performance targets (e.g., the center wavelength accuracy, spectral ripple and chromatic dispersion within the passband, spectral crosstalk outside the passband). As hyperscale data centers are driving radical changes because of the

complexity and integration of new devices and new applications for higher baud rate, higher-order modulations, and higher channel density, some highly desirable attributes of CWDM Mux/Demux components are:

1. Athermal operation. Stable performance including stable wavelength transmission spectra over the operating temperature range removes need for active control and any associated tuning power.
2. Standard 20-nm channel spacing over O-band CWDM wavelengths. This enables uncooled laser operation and will benefit from lower fiber loss and the lowest dispersion in the single mode fiber. One of the earlier 4×25 Gbps 100 G LR4 WDM products employed a 4.5-nm channel spacing to reduce the 10-km fiber dispersion, but it demanded semicooled laser transmitters, more power, and higher cost.
3. Flat-top, clear channel passband of 13 nm and low dispersion within the passband. This passband is wide enough to accommodate the uncooled laser and is agnostic to different modulation rates and formats. Passband ripple and dispersion have not been the main concern for short distance applications until now, but as the data center moves to a higher baud rate (from 25 Gbaud in the past to 53 Gbaud now to 100 Gbaud in the future), minimizing ripple, phase induced dispersion, and inherent higher-order dispersion-induced distortion within the passband is desirable.
4. Low optical crosstalk. This must be significantly improved for PAM4 modulation and high data rate, compared to NRZ signaling, to minimize the power penalty for the same BER at the receiver. Low crosstalk demands a square filter shape, steep slope, and minimal spectral ripples within the passband.
5. Low insertion loss, low polarization-dependent loss and low polarization wavelength dependence. Low loss within the passband translates to improved transmitter OMA and receiver sensitivity and is especially important as the data center is evolving from 100-G NRZ systems to 400 /800-G PAM4 systems.
6. Low optical back-reflection for integrated Mux and Demux in passive photonic integrated chips. It is desired to keep on-chip back-reflection low to minimize the unwanted signal integrity issues, such as multipath interferences and laser feedback tolerance.
7. More integrated Mux/Demux. Advancements in passive photonics integrated circuits can ideally serve the platform for multichannel integration and compact chip designs, with the goal of reducing component cost.

Optical Mux/Demux and filter are divided into two categories: finite impulse response (FIR) optical filter and infinite impulse response (IIR) optical filter, similar to the electric digital filter [122]. FIR optical filters are commonly implemented in feed-forward grating or interferometers in passive photonic circuits, for example, array waveguide gratings (AWGs), echelle grating (EG) and cascaded asymmetric Mach−Zehnder interferometers. The FIR optical filter has linear phase response and low chromatic dispersion within the passband, can easily multiplex/demultiplex multichannel wavelengths, and is less sensitive to process variations, but it has a tradeoff of larger chip size due to the FIR nature. Athermal operating AWGs and EGs, based on a single grating structure, are commercially deployed in the datacom (see Section 5.3.5).

IIR optical filters can be designed with both transmissive and feedback spectral responses in passive photonic circuits, such as microring resonant and contra-DC with Bragg grating (CDC). The IIR optical filter selects one wavelength at each IIR

stage, so it typically needs many devices in series for multiwavelength applications. Its phase response is nonlinear as a function of optical frequency, so careful design is required to minimize the chromatic dispersion inside the filter passband [122,123]. It has a compact device size due to its steep response but is, as a result, also very sensitive to process variations. So, it requires thermal tuning to compensate for process induced deviations in center wavelength and passband shape. These IIR filters have been demonstrated (see Section 5.3.3) but have not yet proven manufacturability, low cost, full dynamic control in actual systems for product adoption. Another IIR filter is the discrete thin-film multilayer cavity filter in commercial deployment. It is not an integrated device so is not included for further discussion.

Table 5.5 summarizes a survey of key Mux/Demux/filter device types, design principles, key performance metrics, impact of manufacturing variations, operating conditions on performance, and state-of-the-art research, development and deployment. Detailed discussion is in the following sections.

5.3.2 Interferometer filters and Mux/Demux

Wave interference is at the core of photonics and is the basic optical principle that enables wavelength filtering, multiplexing and demultiplexing. Specifically, when two waveguides are combined together, light in the two paths will either constructively or destructively interfere, based on the relative phase difference (destructive interference occurs when there is a 180 degrees or, equivalently, π phase shift). The most common approach for realizing interference is with the use of a MZI, which first splits light into two waveguide paths, each of a different length, and then combines them back together. Since the two paths are different lengths, the MZI will have a sinusoidally oscillating transmission, with (ideally) 100% transmission for wavelengths that have waves that constructively interfere when they are combined and (ideally) 0% transmission for wavelengths that destructively interfere, and between 0% and 100% transmission for all of the wavelengths in between that only partially interfere. The repetition between wavelengths that constructively (or destructively) interfere is known as the free spectral range (FSR) and is inversely proportional to the length difference between the two paths, ΔL:

$$\mathrm{FSR}_\lambda = \frac{\lambda^2}{\Delta L\, n_g}$$

where n_g is the group index of the waveguides (assuming both are of identical widths/heights and materials). This asymmetric (also known as unbalanced) MZI repetitively filters a wavelength spectrum with a sinusoidal transmission curve that has a bandwidth that is determined by the path length difference (ΔL). For typical silicon strip waveguides (500 nm $\times$ 220 nm), the FSR will be approximately $\mathrm{FSR}_\lambda \approx 5.5$ nm with $\Delta L = 100$ μm at C + L band wavelengths. As a result, narrower filtering requires longer $\Delta L's$ that approach hundreds of microns.

Table 5.5 Multiplexers/demultiplexers and filters.

	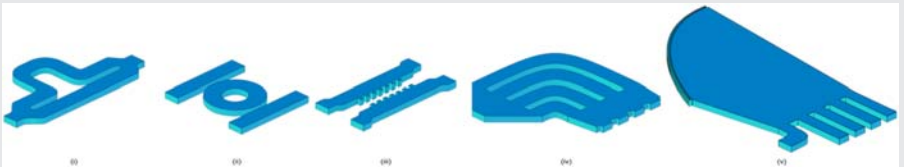				
	(i) Interferometer	**(ii) Ring resonators**	**(iii) Bragg gratings**	**(iv) Arrayed waveguide gratings**	**(v) Echelle gratings**
Dense WDM	No	Yes	No	Yes	No
Coarse WDM	Yes	No	Yes	Yes	Yes
Footprint	Large (mm^2)	Small (μm^2)	1000 s μm^2	Micron × 100s of micron	Micron × tens of microns
Insertion loss	Low (<1.5 dB)	Low (<1 dB)	<1 dB	High	High
Bandwidth	Nanometers	Picometers	Nanometers	Picometers−nanomaters	Picometers−nanomaters
Fundamental transmission profile	Sinusoidal	Lorentzian	Sync	Gaussian	Gaussian
Transmission band flatness	Easy to medium	Difficult	Easy	Difficult	Difficult
Challenges	Thermal sensitivity	Fabrication and thermal sensitivity	Manufacturing complexity and variability, thermal sensitivity	Insertion loss	Insertion loss
In-band extinction ratio (dip of the notch)	Low	Medium	High	N/A	N/A
Channel crosstalk (out of band extinction)	High	Medium	Low	Medium to low	Medium to low
Commercial maturity	Mature (for some applications)	Early	Very mature (lasers), in development (filters)	Very mature	Mature

The basic unbalanced MZI can be used for wavelength multiplexing/demultiplexing by replacing the splitter/combiner with four-port couplers, such as directional or adiabatic couplers [11]. The additional ports can be used to either add or remove wavelengths. By cascading multiple asymmetric MZIs with different FSRs, it is possible to arbitrarily combine/separate wavelength channels.

The filtering performance of an MZI is primarily determined by the quality of the couplers. Perfect destructive interference can only be achieved when the two waves that are being combined have exactly the same amplitude and are exactly a half-wavelength (π) out of phase. Consequently, the couplers must have 50:50 splitting/combining ratios to ensure that both paths of the MZI have the same amplitude. However, DCs have a sinusoidally oscillating coupling ratio that, to the first order, only achieves 50:50 at a single wavelength. Consequently, alternative designs that achieve broadband operation [124–127], such as adiabatic [87] or MMI couplers [63], are typically used. These designs come with tradeoffs in insertion loss, bandwidth, or footprint.

A flat-top spectral response is desired for both filtering and multiplexing applications. The sinusoidal spectral response can be "squared off" with the use of cascaded MZIs to implement an FIR filter response [128]. These multiorder MZI filters are typically known as optical lattice filters [129,130] and achieve a flatter passbands and steep roll-offs by adding additional filtering stages, each of which may consist of multiple MZIs. However, due to fabrication variations, thermal tuning is typically required to realize the desired filter response—which can become complicated with multiple MZI phase shifters and couplers.

It is important to achieve athermal operation so that the filter/Mux operates consistently as the temperature of the chip varies. Silicon has a significant thermo-optic coefficient which results in a varying phase in the MZI arms with temperature. To counteract this, a variety of approaches have been taken, including the use of different width waveguides in each arm of the MZI [131–135], SiN waveguides [136,137] that have a lower thermo-optic coefficient than silicon, and engineered subwavelength metamaterial waveguides [138,139] to optimize the thermo-optic response. Table 5.6 summarizes different demonstrations of MZI Mux/Demux.

5.3.3 Resonator Mux/Demux

High bandwidth optical interconnect in silicon platform can be realized using wavelength division multiplexing (WDM) links, and microring resonator (MR) is one of the fundamental components for achieving these WDM links [10]. With two ports, an input and output (commonly called the through port) ring resonator can be used to filter out signals. For WDM systems an additional waveguide is used to create a four-port device with input, through, drop and add ports. Drop and add ports allow signals to be dropped and added. The maximum # of channels is determined by the wavelength spacing and the ring FSR, which is determined by the ring radius. Bandwidth of the resonance is known as the quality factor and is determined by waveguide-ring coupling strength and the intrinsic loss of the resonator. MRs have shown a promising potential for the photonic WDM multiplexing and demultiplexing

Table 5.6 Mach−Zehnder interferometer Mux/Demux demonstrations.

Design principle	Footprint	Channel number	Channel spacing (nm)	Crosstalk (dB)	Excess loss (dB)	References
Adiabatic splitters	>1 mm	4	3.2	25	2	[140]
Silicon nitride with low thermal sensitivity	1 mm $\times$ 0.6 mm	4	4.5	15−22	<2	[136]
Asymmetric waveguides for a thermal operation	1 mm	4	6.4 nm	15	4	[134]
Polarization insensitive with polarizator rotators, bent directional couplers	1 mm	4	20 nm (CWDM)	>17	<0.25	[141]
Bent directional couplers	0.3 mm $\times$ 0.1 mm	4	20 (CWDM)	20	<1	[126]
Multistage delayed	0.5 mm $\times$ 0.35 mm	4	6.4 nm, 10 nm, 16 nm	24	<2	[142,143]
Point-symmetric wavelength insensitive couplers	<0.5 mm	8	9	10−15	<0.9	[127]
MMI, with thermal tuning	>1 mm	8	3.2	10−20	3.2	[144]
Lattice filters	0.5 mm $\times$ 0.4 mm	8	3.2	20	<1.6	[145]
Ring resonators loaded MZIs, postfab trimming	>1 mm	8	0.133 nm (DWDM)	22	>1	[146]
Automatic phase error correction using heaters	3.3 mm $\times$ 1.6 mm	16	50 GHz	26	2.8−4.1	[147]

system owing to their small footprint and low energy of operation. MR is also used to realize compact modulators in WDM systems and can require a smaller drive voltage compared to other types of photonic modulators, such as Mach–Zehnder modulators (see Chapter 3). Due to the high wavelength selectivity of MRs, many MRs can be cascaded on a single bus waveguide for multichannel WDM architecture to modulate and filter individual channels without the need for a separate Mux/Demux stage as shown in Fig. 5.3 [10].

Several research groups have demonstrated MR-based Mux/Demux in literature. The first ring resonator silicon photonic WDM was demonstrated by Michal Lipson in 2006 [148]. The four-ring resonators had slightly different radii (4.98, 5.00, 5.02, 5.04 μm) to realize a channel spacing of 1.3 nm for the WDM. Heaters are required to fine-tune the resonant frequencies due to fabrication imperfections. The channel spacing can also be controlled by changing the waveguide width in each ring resonator, as was done in [149] to realize a 100-GHz (0.5 nm) WDM system with four channels. Table 5.7 shows a summary for the key MR-based WDM demonstrations so far. It is noted that the key concern of the MR-based WDM system is the operating temperature stability because the operation of the MR-based WDM system fully depends on the wavelength resonance principle of the ring resonator, and the resonance wavelength is very sensitive to the environment temperature. So, to achieve a high speed and stable MR-based WDM Mux/Demux system, temperature control or the integration of temperature insensitive materials/designs is needed.

5.3.4 Waveguide Bragg grating Mux/Demux and filters

Waveguide Bragg gratings are a key component for many integrated photonic devices including lasers, where they function as mirrors, wavelength selective filters and Mux [157]. A single waveguide Bragg grating, being a two-port device, has a strong reflection at its central wavelength (Bragg wavelength), and the transmission response is a band reject filter. Two such waveguide Bragg gratings, when placed in close proximity similar to a DC, can be used to implement a four-port device, which functions as an add-drop Mux. The waveguides are designed to be asymmetric and the grating provides contra-directional coupling between the two waveguides. The concept was first proposed by Elachi and Yeh in 1973 [158] and became practical with high-resolution silicon photonics fabrication [159]. Silicon photonics allows for large grating corrugations and close proximity between waveguides for strong coupling; these two factors give rise to the freedom to design the

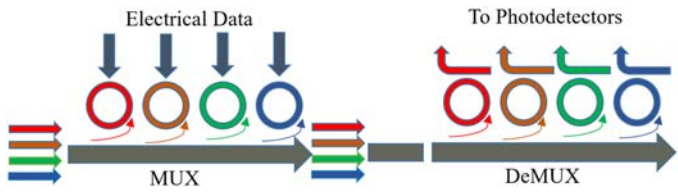

Figure 5.3 A block diagram of microring-based WDM Mux/Demux interconnect.

Table 5.7 Microring WDM Mux/Demux demonstrations.

Design principle	Ring radius (μm)	Channel number	Channel spacing (nm)	Bandwidth/ channel (Gb/s)	Total Bandwidth (Gb/s)	Crosstalk (dB)	Insertion loss (dB)	References
Waveguide width variation and thermally tuned	5	24	0.28	22	528	–	–	[150]
Radius variation	4.98	4	1.3–3.6	4	16	0	0.1	[148]
Waveguide width variation	13	4	0.5	10	40	7	1	[149]
Radius variation Heaters	5	8	1	–	–	0.03	3.5	[151]
Thermally tuned	12	4	1.6	10	40	0	1	[152]
Radius variation	5	5		10	50	–	–	[153]
Radius variation	5	4	3	12.5	50	–	–	[154]
Radius variation	7.5	8	1.6	14	112	–	–	[155]
Thermally tuned	–	16	1	25	400	–	–	[156]

bandwidth of the filter in the range of sub-nm to tens of nanometers. The period of the grating determines the center wavelength of the filter. The filter has an intrinsic flat-band response. Apodization, where the strength of the grating is varied along the length of the device, is used for reducing the side lobes in the spectral response.

Contra-directional Bragg grating filters can be cascaded to improve on their performance. Multiple series-cascaded filters can enhance the extinction ratio of the band-reject filter response (through port). Similarly it can improve the out-of-band rejection ratio of the pass-band filter response (drop port). Rejection ratios of $>$50 dB have been demonstrated. High extinction ratio filters are useful in quantum photonics applications, such as rejecting the laser pump from the output of nonlinear optical devices like single photon sources [160].

Multiple filters with different center wavelengths can be cascaded to build a multichannel filter, and the bandwidth and channel spacing can be engineered for different applications, such as CWDM (20-nm spacing) or LWDM (6.4-nm spacing) [161].

The bandwidth of the filter can be made very large, by introducing a chirp in the grating period along the length of the device, or by using subwavelength waveguides to increase the grating coupling coefficient [157]. Bragg-based filters with subwavelength engineering are compatible with DUV foundry fabrication [162]. For broadband applications, it is also important to eliminate back-reflections from the Bragg gratings, which can be done using an antireflection design where two gratings are out of phase [163].

5.3.5 Arrayed waveguide grating and echelle diffraction gratings Mux/Demux

AWGs and echelle diffraction gratings (EDGs) are commonly used for integrated photonic Mux/Demux and especially in datacom where WDM is used to increase data bandwidth. Both AWG and EDG use multibeam interference to combine spectrally varying signals spatially from 1 to N waveguides, where N can be from tens to hundreds (see Table 5.8) [164,165]. The basic operation of an AWG consists of input/output waveguides which typically lead to areas of free space propagation and then to N number of arrayed waveguides. Unlike AWG, EDG do not have arrayed waveguides but uses an input waveguide which leads to an area of free space propagation and then to a grating, the output waveguides being spaced to collect spectral separated light spatially. Thus AWG and EDG achieve spatial multiplexing by utilizing dispersion, AWG by waveguides of varying length and EDG by a path length difference resulting from grating teeth position [164].

AWG for WDM in telecommunications has been dominated by silica [164,165], but for applications where size and cost rather than pure performance dominate, material platforms, such as Si, SiN, and InP, have been demonstrated with device footprints many orders of magnitude smaller. There are a number of approaches to reduce AWG footprint. One popular strategy is centered on using the large free space propagation areas for both incoming and outgoing signals. Another reuses the

Table 5.8 Arrayed waveguide grating and echelle diffraction grating Mux/Demux.

Type	Area (mm^2)	Channel number	Channel spacing (nm)	Crosstalk (dB)	Excess loss (dB)	References
AWG Si-strip	0.002−0.185	5−50	1.6−11	− 5 to −25	1.1 to 10	[164,166−172]
AWG Si-rib	0.1−176	32−512	0.2−1.6	− 4 to −18	8	[169,171]
AWG SiN-strip	1.17	32	1.6	− 13	1.8	[173]
AWG SiO_2-strip	1200	64−400	0.2−0.8	− 25 to −32	3.4	[171,174,175]
EDG Si-strip	0.016−0.5	4−30	1−20	− 10 to −35	− 1.5 to 10	[164,176−181]
EDG SiN-strip	0.188	4	20	–	–	[182]

free propagation region by looping of the arrayed waveguides around so input/output beams cross perpendicularly within the free space propagation region [166]. Yet, another uses a folded geometry with photonic crystal retro-reflectors, where input/output beams travel in parallel but bi-directionally [183]. Furthermore, a serpentine path for the arrayed waveguides section of AWG has been used to increase path length difference while reducing device size [164], and notably, MZI interlevers have been used to double channel number and halving channel spacing while maintaining a small footprint [184]. To reduce footprint in EDG, Bragg grating and sidewall grating based reflection have been utilized at the end of the free propagation region rather than the traditional grating structure [176–178].

Smaller footprints are driven by the higher refractive index contrast of these waveguide platforms. However, the higher refractive index contrast that allows or these reductions in footprint also create more stringent requirement on dimensional and refractive index control. In both AWG and EDG, fabrication imperfections will cause optical phase errors—manifesting as increased cross talk between channels [168]. Photonics designers can typically expect $\sim \pm 1$ nm of thickness variation over a device area of 1 mm^2 thus must build these considerations into models for AWG and EDG designs [168].

5.4 Future directions

Future photonic chips for datacom communication applications will require continued advancements, including, extra bands for more bandwidth, smaller footprints, lower insertion loss, and lower crosstalk. This is particularly crucial for the design and implementation of future passive devices and circuits to continue improving in performance, cost, size, functionality, and integration. Specifically, the following general trends will be required:

- Lower loss. It is always a fundamental area of focus for reducing waveguide propagation loss and component loss to enable a better power link budget.
- More compact passive device size. The complex silicon photonics PIC involves hundreds and thousands of passive devices for routing, splitting, coupling, interfering, and mode shaping. The shrinking of each passive device is essential to reduce the PIC die size and PIC die cost.
- More complex passive functionality. An example is that polarization-diversity passives photonics circuits will add more complex functionality for potential coherent or polarization-independent chip applications. It will also be used to enable more advanced multiplexing techniques that combine multiple degrees of freedom of light.
- Tunability. The benefit of tunability of passive devices is that allows for the correction of process variations and the reduction of inventory and component numbers. However, tunability comes with tradeoffs that remain hurdles that need to be overcome. Specifically, tunability needs extra complex IC control circuitry at the line or client modules. This brings up the higher-level system questions, such as control feedback accuracy versus system performance, extra power dissipation, manufacturing, reliability, and cost. Ultimately,

it is not clear if the tunability feature will eventually be implemented in systems. There is more research needed in this area, especially for new application explorations.
- Trimming is an alternative to permanently adjusting the material refractive index, compensating for the process variability, and potentially removing the thermal tuning. More research is needed in this area for practical use.

5.4.1 Future directions in design and testing

To facilitate these advancements, there needs to be continued improvements in design and testing methodologies. In terms of advanced passive design approaches, the practice of using inverse design and other machine learning techniques is quickly emerging as a key tool for device optimization. For example, an ultra-compact 4-ch CWDM Demux was reported using inverse design [185], although the performance and demonstration are still in an early stage.

In terms of performance, with the increase in foundry scale integrated photonics, we are now at a point where measuring statistical variation has become a critical need. Specifically, predicting the spread in circuit performance in the presence of manufacturing variations provides valuable information about expected functional yield and the stability of the fabrication process, both of which are vital to determining the total cost of the PICs. The underlying manufacturing-related variations arise from fluctuations in process conditions and manifest in passive devices as spread in waveguide dimensions [11,186,187] across dies, wafers, and lots. For passive devices, the process variations result in effective index changes, which shift the circuit characteristics. Active devices have additional variations in doping, contact and metal resistances. In addition to manufacturing variations, advanced effects, such as correlated variation in output characteristics between closely spaced (inadvertently coupled) components [11,187] and layout dependent effects, need to be considered.

Optimization of passive device components will eventually yield devices with losses so low that it will become challenging to measure them. Therefore an open question is how does one test components with extremely small (milli-dB) losses? Related to the previous points on statistical variations, measuring inadvertent coupling (layout dependent) between components is also an outstanding challenge due to the coherent nature of light. New imaging techniques and test circuits need to be developed to characterize future PICs with the dynamic range required for demanding applications, such as advanced multiplexing, DWDM, quantum, and more.

5.4.2 Future directions in integration

While silicon photonics has made extraordinary advancements that have led to a photonic datacom revolution, there are significant tradeoffs/challenges to a CMOS-only photonic platform. The challenge of laser integration still remains, and free-carrier plasma dispersion-based modulators still have significant tradeoffs in performance (V_π, energy, footprint) versus loss. Consequently, the overall silicon photonics field continues to explore the heterogeneous integration of materials to advance on these

shortfalls. For example, lithium niobate [188], electro-optic polymers [189], and BTO [190] are being integrated to achieve high performance Pockels-based electro-optic modulation. Similarly, III−V gain material [6] is being integrated for lasers and amplifiers; aluminum nitride [191] for piezo tuning, high power handling, and ultra-broad operating bandwidth into the SWIR-UV; and micro-electro-mechanical for low power phase tuning [192]. Passive devices must be developed to achieve low loss transitions between all of these heterogenous materials. Lastly, just as the past several decades in datacom advancements have pushed the PIC platform forward, we expect that these research areas to enable many future applications (Terabit Datacom, Artificial Intelligence, Quantum Information, Sensing, LIDAR, and more).

References

[1] Q. Cheng, M. Bahadori, M. Glick, S. Rumley, K. Bergman, Recent advances in optical technologies for data centers: a review, Optica 5 (11) (2018) 1354−1370. Available from: https://doi.org/10.1364/OPTICA.5.001354.

[2] P.P. Absil, et al., Imec iSiPP25G silicon photonics: a robust CMOS-based photonics technology platform, San Francisco, California, United States, February 2015, p. 93670V. https://doi.org/10.1117/12.2076262.

[3] K. Giewont, et al., 300-mm monolithic silicon photonics foundry technology, IEEE J. Sel. Top. Quantum Electron. 25 (5) (2019) 1−11. Available from: https://doi.org/10.1109/JSTQE.2019.2908790.

[4] N.M. Fahrenkopf, C. McDonough, G.L. Leake, Z. Su, E. Timurdogan, D.D. Coolbaugh, The AIM photonics MPW: a highly accessible cutting edge technology for rapid prototyping of photonic integrated circuits, IEEE J. Sel. Top. Quantum Electron. 25 (5) (2019). Available from: https://doi.org/10.1109/JSTQE.2019.2935698.

[5] S.Y. Siew, et al., Review of silicon photonics technology and platform development, J. Light. Technol. 39 (13) (2021) 4374−4389. Available from: https://doi.org/10.1109/JLT.2021.3066203.

[6] J.C. Norman, D. Jung, Y. Wan, J.E. Bowers, Perspective: the future of quantum dot photonic integrated circuits, APL Photonics 3 (3) (2018) 030901. Available from: https://doi.org/10.1063/1.5021345.

[7] L. Carpenter, et al., Towards low propagation losses in active photonic multi-project wafer runs, 2021.

[8] Zettabyte Era. <https://en.wikipedia.org/wiki/Zettabyte_Era>.

[9] G. Keeler, Photonics in the package for extreme scalability (PIPES). <https://www.darpa.mil/program/photonics-in-the-package-for-extreme-scalability>.

[10] Y. London, et al., Performance requirements for terabit-class silicon photonic links based on cascaded microring resonators, J. Light. Technol. 38 (13) (2020) 3469−3477. Available from: https://doi.org/10.1109/JLT.2019.2961455.

[11] L. Chrostowski, M. Hochberg, Silicon Photonics Design: From Devices to Systems, Cambridge University Press, Cambridge, 2015. Available from: https://doi.org/10.1017/CBO9781316084168.

[12] S.K. Selvaraja, et al., Highly uniform and low-loss passive silicon photonics devices using a 300 mm CMOS platform, in: OFC 2014, 2014, pp. 1−3. https://doi.org/10.1364/OFC.2014.Th2A.33.

[13] Q. Wilmart, et al., A complete Si photonics platform embedding ultra-low loss waveguides for O- and C-band, J. Light. Technol. 39 (2) (2021) 532−538. Available from: https://doi.org/10.1109/JLT.2020.3030123.
[14] M.A. Tran, D. Huang, T. Komljenovic, J. Peters, A. Malik, J.E. Bowers, Ultra-low-loss silicon waveguides for heterogeneously integrated silicon/III−V photonics, Appl. Sci. 8 (7) (2018). Available from: https://doi.org/10.3390/app8071139.
[15] P. Muñoz, et al., Silicon nitride photonic integration platforms for visible, near-infrared and mid-infrared applications, Sensors 17 (9) (2017). Available from: https://doi.org/10.3390/s17092088.
[16] C.G.H. Roeloffzen, et al., Low-loss Si3N4 TriPleX optical waveguides: technology and applications overview, IEEE J. Sel. Top. Quantum Electron. 24 (4) (2018) 1−21. Available from: https://doi.org/10.1109/JSTQE.2018.2793945.
[17] J. Liu, et al., High-yield, wafer-scale fabrication of ultralow-loss, dispersion-engineered silicon nitride photonic circuits, Nat. Commun. 12 (1) (2021). Available from: https://doi.org/10.1038/s41467-021-21973-z.
[18] G.E. Hoefler, et al., Foundry development of system-on-chip InP-based photonic integrated circuits, IEEE J. Sel. Top. Quantum Electron. 25 (5) (2019) 1−17. Available from: https://doi.org/10.1109/JSTQE.2019.2906270.
[19] C. Ciminelli, F. Dell'Olio, M.N. Armenise, F.M. Soares, W. Passenberg, High performance InP ring resonator for new generation monolithically integrated optical gyroscopes, Opt. Express 21 (1) (2013) 556−564. Available from: https://doi.org/10.1364/OE.21.000556.
[20] J.J.G.M. van der Tol, et al., Indium phosphide integrated photonics in membranes, IEEE J. Sel. Top. Quantum Electron. 24 (1) (2018) 1−9. Available from: https://doi.org/10.1109/JSTQE.2017.2772786.
[21] J. Wang, et al., Gallium arsenide (GaAs) quantum photonic waveguide circuits, Opt. Commun. 327 (2014) 49−55. Available from: https://doi.org/10.1016/j.optcom.2014.02.040.
[22] K. Okamoto, Planar lightwave circuits, in: K. Okamoto (Ed.), Fundamentals of Optical Waveguides, second ed., Academic Press, Burlington, 2006, pp. 417−534. Available from: https://doi.org/10.1016/B978-012525096-2/50010-6.
[23] M. Smit, K. Williams, J. van der Tol, Past, present, and future of InP-based photonic integration, APL Photonics 4 (5) (May 2019) 050901. Available from: https://doi.org/10.1063/1.5087862.
[24] W. Bogaerts, et al., Programmable photonic circuits, Nature 586 (7828) (2020) 207−216.
[25] T. Aalto, et al., Open-access 3-μm SOI waveguide platform for dense photonic integrated circuits, IEEE J. Sel. Top. Quantum Electron. 25 (5) (2019) 1−9. Available from: https://doi.org/10.1109/JSTQE.2019.2908551.
[26] A.L. Gaeta, M. Lipson, T.J. Kippenberg, Photonic-chip-based frequency combs, Nat. Photonics 13 (3) (2019). Available from: https://doi.org/10.1038/s41566-019-0358-x.
[27] S. Chung, H. Abediasl, H. Hashemi, A monolithically integrated large-scale optical phased array in silicon-on-insulator CMOS, IEEE J. Solid State Circuits 53 (1) (2018) 275−296. Available from: https://doi.org/10.1109/JSSC.2017.2757009.
[28] A. Dutt, K. Luke, S. Manipatruni, A.L. Gaeta, P. Nussenzveig, M. Lipson, On-chip optical squeezing, Phys. Rev. Appl. 3 (4) (2015) 044005. Available from: https://doi.org/10.1103/PhysRevApplied.3.044005.
[29] M.H.P. Pfeiffer, J. Liu, A.S. Raja, T. Morais, B. Ghadiani, T.J. Kippenberg, Ultra-smooth silicon nitride waveguides based on the Damascene reflow process: fabrication and loss origins, Optica 5 (7) (2018) 884−892. Available from: https://doi.org/10.1364/OPTICA.5.000884.

[30] W. Jin, D.D. John, J.F. Bauters, T. Bosch, B.J. Thibeault, J.E. Bowers, Deuterated silicon dioxide for heterogeneous integration of ultra-low-loss waveguides, Opt. Lett. 45 (12) (2020) 3340−3343. Available from: https://doi.org/10.1364/OL.394121.
[31] A.H. Atabaki, et al., Integrating photonics with silicon nanoelectronics for the next generation of systems on a chip, Nature 556 (7701) (2018). Available from: https://doi.org/10.1038/s41586-018-0028-z.
[32] N.F. Tyndall, et al., A low-loss SiN photonic integrated circuit foundry platform for waveguide-enhanced Raman spectroscopy, Smart Photonic Optoelectronic Integr. Circuits XXIII, 11690, 2021, pp. 40−46. Available from: http://doi.org/10.1117/12.2582529.
[33] J.H. Song, T.D. Kongnyuy, N. Hosseini, A. Stassen, R. Jansen, X. Rottenberg, Advanced waveguide bends for photonic integrated circuits, in: 45th European Conference on Optical Communication (ECOC 2019), September 2019, pp. 1−3. https://doi.org/10.1049/cp.2019.1087.
[34] G.T. Reed, A.P. Knights, Silicon Photonics: An Introduction, John Wiley & Sons, 2004.
[35] T.Y.L. Ang, et al., Versatile Bezier bends for silicon photonics, in: 2017 Conference on Lasers and Electro-Optics Pacific Rim (2017), Paper s1869, July 2017, p. s1869. <https://opg.optica.org/abstract.cfm?uri = CLEOPR-2017-s1869>.
[36] T. Fujisawa, S. Makino, T. Sato, K. Saitoh, Low-loss, compact, and fabrication-tolerant Si-wire 90° waveguide bend using clothoid and normal curves for large scale photonic integrated circuits, Opt. Express 25 (8) (2017) 9150−9159. Available from: https://doi.org/10.1364/OE.25.009150.
[37] Y. Liu, et al., Very sharp adiabatic bends based on an inverse design, Opt. Lett. 43 (11) (2018) 2482−2485. Available from: https://doi.org/10.1364/OL.43.002482.
[38] S. Jahani, et al., Controlling evanescent waves using silicon photonic all-dielectric metamaterials for dense integration, Nat. Commun. 9 (1) (May 2018). Available from: https://doi.org/10.1038/s41467-018-04276-8.
[39] M. Johnson, M. Johnson, M.G. Thompson, D. Sahin, Low-loss, low-crosstalk waveguide crossing for scalable integrated silicon photonics applications, Opt. Express 28 (9) (2020) 12498−12507. Available from: https://doi.org/10.1364/OE.381304.
[40] S. Nevlacsil, P. Muellner, M. Sagmeister, J. Kraft, R. Hainberger, Broadband low loss and ultra-low crosstalk waveguide crossings based on a multimode interferometer for 840 nm operation, OSA Contin. 3 (2) (2020) 334−344. Available from: https://doi.org/10.1364/OSAC.381003.
[41] S. Wu, X. Mu, L. Cheng, S. Mao, H.Y. Fu, State-of-the-art and perspectives on silicon waveguide crossings: a review, Micromachines 11 (3) (2020). Available from: https://doi.org/10.3390/mi11030326.
[42] J. Wallner, AIM photonics PDK documentation. <https://www.aimphotonics.com/pdk>.
[43] W. Bogaerts, P. Dumon, D.V. Thourhout, R. Baets, Low-loss, low-cross-talk crossings for silicon-on-insulator nanophotonic waveguides, Opt. Lett. 32 (19) (2007) 2801−2803. Available from: https://doi.org/10.1364/OL.32.002801.
[44] Y. Ma, et al., Ultralow loss single layer submicron silicon waveguide crossing for SOI optical interconnect, Opt. Express 21 (24) (2013) 29374−29382. Available from: https://doi.org/10.1364/OE.21.029374.
[45] T.J. Seok, N. Quack, S. Han, R.S. Muller, M.C. Wu, Large-scale broadband digital silicon photonic switches with vertical adiabatic couplers, Optica 3 (1) (2016) 64−70. Available from: https://doi.org/10.1364/OPTICA.3.000064.
[46] Y. Zhang, A. Hosseini, X. Xu, D. Kwong, R.T. Chen, Ultralow-loss silicon waveguide crossing using Bloch modes in index-engineered cascaded multimode-interference

couplers, Opt. Lett. 38 (18) (2013) 3608−3611. Available from: https://doi.org/10.1364/OL.38.003608.
[47] D. Celo, et al., Optical proximity correction in geometry sensitive silicon photonics waveguide crossings, in: 2017 IEEE 14th International Conference on Group IV Photonics (GFP), August 2017, pp. 45−46. https://doi.org/10.1109/GROUP4.2017.8082188.
[48] P.J. Bock, et al., Subwavelength grating crossings for silicon wire waveguides, Opt. Express 18 (15) (2010) 16146−16155. Available from: https://doi.org/10.1364/OE.18.016146.
[49] Z. Dong, et al., Ultra-compact X-shaped waveguide crossings with flexible angles based on inverse design, Opt. Express 29 (13) (2021) 19715−19726. Available from: https://doi.org/10.1364/OE.426733.
[50] M.M. Gilarlue, J. Nourinia, Ch Ghobadi, S.H. Badri, H. Rasooli Saghai, Multilayered Maxwell's fisheye lens as waveguide crossing, Opt. Commun. 435 (2019) 385−393. Available from: https://doi.org/10.1016/j.optcom.2018.11.057.
[51] B. Wu, Y. Yu, X. Zhang, Multimode waveguide crossing with ultralow loss and low imbalance, Opt. Express 28 (10) (May 2020) 14705−14711. Available from: https://doi.org/10.1364/OE.392445.
[52] W.D. Sacher, et al., Multilayer silicon nitride-on-silicon integrated photonic platform for 3D photonic circuits, in: 2016 Conference on Lasers and Electro-Optics (CLEO), June 2016, pp. 1−2.
[53] K. Shang, S. Pathak, B. Guan, G. Liu, S.J.B. Yoo, Low-loss compact multilayer silicon nitride platform for 3D photonic integrated circuits, Opt. Express 23 (16) (2015) 21334−21342. Available from: https://doi.org/10.1364/OE.23.021334.
[54] E. Timurdogan, et al., APSUNY process design kit (PDKv3.0): O, C and L band silicon photonics component libraries on 300mm wafers, in: Optical Fiber Communication Conference (OFC) 2019 (2019), Paper Tu2A.1, March 2019, p. Tu2A.1. https://doi.org/10.1364/OFC.2019.Tu2A.1.
[55] L.G. Carpenter, et al., Low losses in an active silicon photonic multi-project wafer runs, Silicon Photonics XVII, vol. 12006, 2022, p. 1200602. Available from: http://doi.org/10.1117/12.2609355.
[56] S. Yang, et al., A single adiabatic microring-based laser in 220 nm silicon-on-insulator, Opt. Express 22 (1) (2014) 1172−1180. Available from: https://doi.org/10.1364/OE.22.001172.
[57] J.H. Lee, et al., High power and widely tunable Si hybrid external-cavity laser for power efficient Si photonics WDM links, Opt. Express 22 (7) (2014) 7678−7685. Available from: https://doi.org/10.1364/OE.22.007678.
[58] Y. Zhang, et al., Sagnac loop mirror and micro-ring based laser cavity for silicon-on-insulator, Opt. Express 22 (15) (2014) 17872−17879. Available from: https://doi.org/10.1364/OE.22.017872.
[59] G.-H. Duan, et al., Hybrid III−V on silicon lasers for photonic integrated circuits on silicon, IEEE J. Sel. Top. Quantum Electron. 20 (4) (2014) 158−170. Available from: https://doi.org/10.1109/JSTQE.2013.2296752.
[60] A. Sakai, T. Fukazawa, T. Baba, Low loss ultra-small branches in a silicon photonic wire waveguide, IEICE Trans. Electron. (2002) 6.
[61] Y. Zhang, et al., A compact and low loss Y-junction for submicron silicon waveguide, Opt. Express 21 (1) (2013) 1310−1316. Available from: https://doi.org/10.1364/OE.21.001310.
[62] A.Y. Piggott, et al., Inverse-designed photonics for semiconductor foundries, ACS Photonics 7 (3) (2020) 569−575. Available from: https://doi.org/10.1021/acsphotonics.9b01540.

[63] H. Guan, et al., Compact and low loss 90° optical hybrid on a silicon-on-insulator platform, Opt. Express 25 (23) (2017) 28957. Available from: https://doi.org/10.1364/OE.25.028957.
[64] Y. Wang, et al., Compact broadband directional couplers using subwavelength gratings, IEEE Photonics J. 8 (3) (2016) 1–8. Available from: https://doi.org/10.1109/JPHOT.2016.2574335.
[65] H. Yun, L. Chrostowski, N.A.F. Jaeger, Ultra-broadband 2 × 2 adiabatic 3 dB coupler using subwavelength-grating-assisted silicon-on-insulator strip waveguides, Opt. Lett. 43 (8) (2018) 1935–1938. Available from: https://doi.org/10.1364/OL.43.001935.
[66] J.M.F. Cabanillas, M.K. Singh, B. Zhang, M.A. Popović, Compact broadband rapid-adiabatic polarization splitter-rotators in a monolithic electronic-photonic SOI platform, in: Conference on Lasers and Electro-Optics (2021), Paper SW4E0.2, May 2021, p. SW4E.2. https://doi.org/10.1364/CLEO_SI.2021.SW4E.2.
[67] J.M.F. Cabanillas, et al., Demonstration and fabrication tolerance study of a low-loss, ultra-broadband rapid adiabatic 3-dB coupler in a next-generation 45 nm monolithic electronic-photonic platform, in: Optical Fiber Communication Conference (OFC) 2021 (2021), Paper Tu5B.4, June 2021, p. Tu5B.4. <https://www.osapublishing.org/abstract.cfm?uri = OFC-2021-Tu5B.4>.
[68] Z. Tu, et al., A compact SOI polarization beam splitter based on multimode interference coupler, in: 2011 Asia Communications and Photonics Conference and Exhibition (ACP), November 2011, pp. 1–6. https://doi.org/10.1117/12.905589.
[69] J.M. Hong, et al., Design and fabrication of a significantly shortened multimode interference coupler for polarization splitter application, IEEE Photonics Technol. Lett. 15 (1) (2003) 72–74. Available from: https://doi.org/10.1109/LPT.2002.805803.
[70] X. Sun, M.Z. Alam, J.S. Aitchison, M. Mojahedi, Polarization rotator based on augmented low-index-guiding waveguide on silicon nitride/silicon-on-insulator platform, Opt. Lett. 41 (14) (2016) 3229–3232. Available from: https://doi.org/10.1364/OL.41.003229.
[71] B. Shen, P. Wang, R. Polson, R. Menon, An integrated-nanophotonics polarization beamsplitter with 2.4 × 2.4 μm^2 footprint, Nat. Photonics 9 (6) (2015). Available from: https://doi.org/10.1038/nphoton.2015.80.
[72] D. Dai, Z. Wang, J. Peters, J.E. Bowers, Compact polarization beam splitter using an asymmetrical Mach–Zehnder interferometer based on silicon-on-insulator waveguides, IEEE Photonics Technol. Lett. 24 (8) (2012) 673–675. Available from: https://doi.org/10.1109/LPT.2012.2184530.
[73] T. Yamazaki, H. Aono, J. Yamauchi, H. Nakano, Coupled waveguide polarization splitter with slightly different core widths, J. Light. Technol. 26 (21) (2008) 3528–3533.
[74] D. Dai, J.E. Bowers, Novel ultra-short and ultra-broadband polarization beam splitter based on a bent directional coupler, Opt. Express 19 (19) (2011) 18614–18620. Available from: https://doi.org/10.1364/OE.19.018614.
[75] Y. Fang, et al., Polarization beam splitter based on Si3N4/SiO2 horizontal slot waveguides for on-chip high-power applications, Sensors 20 (10) (2020). Available from: https://doi.org/10.3390/s20102862.
[76] J. Wang, Design of a compact polarization beam splitter for silicon-based cross-slot waveguides, J. Mod. Opt. 67 (15) (2020) 1277–1284. Available from: https://doi.org/10.1080/09500340.2020.1842928.
[77] D. Dai, Y. Tang, J.E. Bowers, Mode conversion in tapered submicron silicon ridge optical waveguides, Opt. Express 20 (12) (2012) 13425–13439. Available from: https://doi.org/10.1364/OE.20.013425.

[78] Y. Ding, L. Liu, C. Peucheret, H. Ou, Fabrication tolerant polarization splitter and rotator based on a tapered directional coupler, Opt. Express 20 (18) (2012) 20021−20027. Available from: https://doi.org/10.1364/OE.20.020021.
[79] D. Dai, J.E. Bowers, Novel concept for ultracompact polarization splitter-rotator based on silicon nanowires, Opt. Express 19 (11) (May 2011) 10940−10949. Available from: https://doi.org/10.1364/OE.19.010940.
[80] J. Zhang, T.-Y. Liow, M. Yu, G.-Q. Lo, D.-L. Kwong, Silicon waveguide based TE mode converter, Opt. Express 18 (24) (2010) 25264−25270. Available from: https://doi.org/10.1364/OE.18.025264.
[81] W. Ye, X. Yuan, Y. Gao, J. Liu, Design of broadband silicon-waveguide mode-order converter and polarization rotator with small footprints, Opt. Express 25 (26) (2017) 33176−33183. Available from: https://doi.org/10.1364/OE.25.033176.
[82] M.R. Watts, H.A. Haus, Integrated mode-evolution-based polarization rotators, Opt. Lett. 30 (2) (2005) 138−140. Available from: https://doi.org/10.1364/OL.30.000138.
[83] L. Chen, C.R. Doerr, Y.-K. Chen, Compact polarization rotator on silicon for polarization-diversified circuits, Opt. Lett. 36 (4) (2011) 469−471. Available from: https://doi.org/10.1364/OL.36.000469.
[84] M.A. Popović, A. Khilo, Integrated photonic magic-T (with twice the magic), Integr. Photonics Research, Silicon Nanophotonics Photonics Switching (2010), Pap. IWG7 (2010) IWG7. Available from: https://doi.org/10.1364/IPRSN.2010.IWG7.
[85] H. Xu, Y. Shi, Ultra-broadband silicon polarization splitter-rotator based on the multimode waveguide, Opt. Express 25 (15) (2017) 18485−18491. Available from: https://doi.org/10.1364/OE.25.018485.
[86] W.D. Sacher, T. Barwicz, B.J.F. Taylor, J.K.S. Poon, Polarization rotator-splitters in standard active silicon photonics platforms, Opt. Express 22 (4) (2014) 3777−3786. Available from: https://doi.org/10.1364/OE.22.003777.
[87] J.M.F. Cabanillas, B. Zhang, M.A. Popović, Demonstration of 3 ± 0.12 dB power splitting over 145 nm optical bandwidth in a 31-μm long 3-dB rapid adiabatic coupler, in: 2020 Optical Fiber Communications Conference and Exhibition (OFC), March 2020, pp. 1−3.
[88] V.R. Almeida, R.R. Panepucci, M. Lipson, Nanotaper for compact mode conversion, Opt. Lett. 28 (15) (2003) 1302. Available from: https://doi.org/10.1364/OL.28.001302.
[89] J. Wang, et al., Low-loss and misalignment-tolerant fiber-to-chip edge coupler based on double-tip inverse tapers, in: Optical Fiber Communication Conference, Anaheim, CA, 2016, p. M2I.6. https://doi.org/10.1364/OFC.2016.M2I.6.
[90] B. Ben Bakir, et al., Low-loss (<1dB) and polarization-insensitive edge fiber couplers fabricated on 200-mm silicon-on-insulator wafers, IEEE Photonics Technol. Lett. 22 (11) (2010) 739−741. Available from: https://doi.org/10.1109/LPT.2010.2044992.
[91] P. Yin, et al., Low connector-to-connector loss through silicon photonic chips using ultra-low loss splicing of SMF-28 to high numerical aperture fibers, Opt. Express 27 (17) (2019) 24188−24193. Available from: https://doi.org/10.1364/OE.27.024188.
[92] T. Barwicz, et al., An O-band metamaterial converter interfacing standard optical fibers to silicon nanophotonic waveguides, in: Optical Fiber Communication Conference, Los Angeles, CA, 2015, p. Th3F.3. https://doi.org/10.1364/OFC.2015.Th3F.3.
[93] M. Pu, L. Liu, H. Ou, K. Yvind, J.M. Hvam, Ultra-low-loss inverted taper coupler for silicon-on-insulator ridge waveguide, Opt. Commun. 283 (19) (2010) 3678−3682. Available from: https://doi.org/10.1016/j.optcom.2010.05.034.
[94] S. McNab, N. Moll, Y. Vlasov, Ultra-low loss photonic integrated circuit with membrane-type photonic crystal waveguides, Opt. Express 11 (22) (2003) 2927. Available from: https://doi.org/10.1364/OE.11.002927.

[95] M. Papes, et al., Fiber-chip edge coupler with large mode size for silicon photonic wire waveguides, Opt. Express 24 (5) (2016) 5026. Available from: https://doi.org/10.1364/OE.24.005026.
[96] R. Halir, et al., Subwavelength-grating metamaterial structures for silicon photonic devices, Proc. IEEE 106 (12) (2018) 2144–2157. Available from: https://doi.org/10.1109/JPROC.2018.2851614.
[97] J. Nauriyal, M. Song, R. Yu, J. Cardenas, Fiber-to-chip fusion splicing for low-loss photonic packaging, Optica 6 (5) (May 2019) 549–552. Available from: https://doi.org/10.1364/OPTICA.6.000549.
[98] M. Wood, P. Sun, R.M. Reano, Compact cantilever couplers for low-loss fiber coupling to silicon photonic integrated circuits, Opt. Express 20 (1) (2012) 164. Available from: https://doi.org/10.1364/OE.20.000164.
[99] J. Cardenas, et al., High coupling efficiency etched facet tapers in silicon, in Conference on Lasers and Electro-Optics 2012 (2012), Paper JW4A.10, May 2012, p. JW4A.10. https://doi.org/10.1364/CLEO_AT.2012.JW4A.10.
[100] S. Malhouitre, et al., Heterogeneous and multi-level integration on mature 25Gb/s silicon photonic platform, in: 2017 IEEE CPMT Symposium Japan (ICSJ), November 2017, pp. 223–226. https://doi.org/10.1109/ICSJ.2017.8240122.
[101] D. Taillaert, et al., An out-of-plane grating coupler for efficient butt-coupling between compact planar waveguides and single-mode fibers, IEEE J. Quantum Electron. 38 (7) (2002) 949–955. Available from: https://doi.org/10.1109/JQE.2002.1017613.
[102] Y. Wang, et al., Focusing sub-wavelength grating couplers with low back reflections for rapid prototyping of silicon photonic circuits, Opt. Express 22 (17) (2014) 20652. Available from: https://doi.org/10.1364/OE.22.020652.
[103] R. Halir, P. Cheben, S. Janz, D.-X. Xu, Í. Molina-Fernández, J.G. Wangüemert-Pérez, Waveguide grating coupler with subwavelength microstructures, Opt. Lett. 34 (9) (May 2009) 1408. Available from: https://doi.org/10.1364/OL.34.001408.
[104] N. Hoppe, et al., Ultra-Efficient Silicon-on-Insulator Grating Couplers With Backside Metal Mirrors, IEEE J. Sel. Top. Quantum Electron. 26 (2) (2020) 1–6. Available from: https://doi.org/10.1109/JSTQE.2019.2935296.
[105] J. Hong, A.M. Spring, F. Qiu, S. Yokoyama, A high efficiency silicon nitride waveguide grating coupler with a multilayer bottom reflector, Sci. Rep. 9 (1) (2019) 12988. Available from: https://doi.org/10.1038/s41598-019-49324-5.
[106] M.T. Wade, et al., 75% efficient wide bandwidth grating couplers in a 45 nm microelectronics CMOS process, in: 2015 IEEE Optical Interconnects Conference (OI), San Diego, CA, April 2015, pp. 46–47. https://doi.org/10.1109/OIC.2015.7115679.
[107] J. Notaros, et al., Ultra-efficient CMOS fiber-to-chip grating couplers, in: Optical Fiber Communication Conference, Anaheim, California, 2016, p. M2I.5. https://doi.org/10.1364/OFC.2016.M2I.5.
[108] C. Alonso-Ramos, P. Cheben, A. Ortega-Moñux, J.H. Schmid, D.-X. Xu, I. Molina-Fernández, Fiber-chip grating coupler based on interleaved trenches with directionality exceeding 95%, Opt. Lett. 39 (18) (2014) 5351. Available from: https://doi.org/10.1364/OL.39.005351.
[109] E.W. Ong, N.M. Fahrenkopf, D.D. Coolbaugh, SiNx bilayer grating coupler for photonic systems, OSA Contin 1 (1) (2018) 13. Available from: https://doi.org/10.1364/OSAC.1.000013.
[110] Z. Zhao, S. Fan, Design principles of apodized grating couplers, J. Light. Technol. 38 (16) (2020) 4435–4446. Available from: https://doi.org/10.1109/JLT.2020.2992574.

[111] L. Cheng, S. Mao, Z. Li, Y. Han, H.Y. Fu, Grating couplers on silicon photonics: design principles, emerging trends and practical issues, Micromachines 11 (7) (2020). Available from: https://doi.org/10.3390/mi11070666.
[112] X. Chen, H.K. Tsang, Polarization-independent grating couplers for silicon-on-insulator nanophotonic waveguides, Opt. Lett. 36 (6) (2011) 796. Available from: https://doi.org/10.1364/OL.36.000796.
[113] C. Alonso-Ramos, L. Zavargo-Peche, A. Ortega-Moñux, R. Halir, I. Molina-Fernández, P. Cheben, Polarization-independent grating coupler for micrometric silicon rib waveguides, Opt. Lett. 37 (17) (2012) 3663. Available from: https://doi.org/10.1364/OL.37.003663.
[114] J.H. Song, F.E. Doany, A.K. Medhin, N. Dupuis, B.G. Lee, F.R. Libsch, Polarization-independent nonuniform grating couplers on silicon-on-insulator, Opt. Lett. 40 (17) (2015) 3941. Available from: https://doi.org/10.1364/OL.40.003941.
[115] J.C.C. Mak, W.D. Sacher, H. Ying, X. Luo, P.G.-Q. Lo, J.K.S. Poon, Multi-layer silicon nitride-on-silicon polarization-independent grating couplers, Opt. Express 26 (23) (2018) 30623−30633. Available from: https://doi.org/10.1364/OE.26.030623.
[116] T. Barwicz, et al., A novel approach to photonic packaging leveraging existing high-throughput microelectronic facilities, IEEE J. Sel. Top. Quantum Electron. 22 (6) (2016) 455−466. Available from: https://doi.org/10.1109/JSTQE.2016.2593637.
[117] T. Barwicz, et al., Advances in interfacing optical fibers to nanophotonic waveguides via mechanically compliant polymer waveguides, IEEE J. Sel. Top. Quantum Electron. 26 (2) (2020) 1−12. Available from: https://doi.org/10.1109/JSTQE.2020.2964383.
[118] R. Dangel, et al., Polymer waveguides enabling scalable low-loss adiabatic optical coupling for silicon photonics, IEEE J. Sel. Top. Quantum Electron. 24 (4) (2018) 1−11. Available from: https://doi.org/10.1109/JSTQE.2018.2812603.
[119] S. Khan, S.M. Buckley, J. Chiles, R.P. Mirin, S.W. Nam, J.M. Shainline, Low-loss, high-bandwidth fiber-to-chip coupling using capped adiabatic tapered fibers, APL Photonics 5 (5) (May 2020) 056101. Available from: https://doi.org/10.1063/1.5145105.
[120] N. Lindenmann, et al., Photonic wire bonding: a novel concept for chip-scale interconnects, Opt. Express 20 (16) (2012) 17667. Available from: https://doi.org/10.1364/OE.20.017667.
[121] N. Lindenmann, et al., Connecting silicon photonic circuits to multicore fibers by photonic wire bonding, J. Light. Technol. 33 (4) (2015) 755−760. Available from: https://doi.org/10.1109/JLT.2014.2373051.
[122] C.K. Madsen, J.H. Zhao, Optical Filter Design and Analysis, Wiley, New York, 1999.
[123] J.R. Ong, R. Kumar, S. Mookherjea, Ultra-high-contrast and tunable-bandwidth filter using cascaded high-order silicon microring filters, IEEE Photonics Technol. Lett. 25 (16) (2013) 1543−1546. Available from: https://doi.org/10.1109/LPT.2013.2267539.
[124] Z. Lu, et al., Broadband silicon photonic directional coupler using asymmetric-waveguide based phase control, Opt. Express 23 (3) (2015) 3795−3808. Available from: https://doi.org/10.1364/OE.23.003795.
[125] J. Van Campenhout, W.M.J. Green, S. Assefa, Y.A. Vlasov, Low-power, 2 x 2 silicon electro-optic switch with 110-nm bandwidth for broadband reconfigurable optical networks, Opt. Express 17 (26) (2009) 24020−24029. Available from: https://doi.org/10.1364/OE.17.024020.
[126] H. Xu, Y. Shi, Flat-top CWDM (de)multiplexer based on MZI with bent directional couplers, IEEE Photonics Technol. Lett. 30 (2) (2018) 169−172. Available from: https://doi.org/10.1109/LPT.2017.2779489.

[127] S.-H. Jeong, Broadband 1 × 8 channel silicon-nanowire-waveguide WDM filter based on point-symmetric Mach-Zehnder interferometric optical couplers in the O-band spectral regime, OSA Contin. 2 (12) (2019) 3564. Available from: https://doi.org/10.1364/OSAC.2.003564.
[128] K. Jinguji, M. Oguma, Optical half-band filters, J. Light. Technol. 18 (2) (2000) 252–259. Available from: https://doi.org/10.1109/50.822800.
[129] S. Ibrahim, et al., Demonstration of a fast-reconfigurable silicon CMOS optical lattice filter, Opt. Express 19 (14) (2011) 13245–13256. Available from: https://doi.org/10.1364/OE.19.013245.
[130] L. Lu, et al., CMOS-compatible temperature-independent tunable silicon optical lattice filters, Opt. Express 21 (8) (2013) 9447–9456. Available from: https://doi.org/10.1364/OE.21.009447.
[131] B. Guha, A. Gondarenko, M. Lipson, Minimizing temperature sensitivity of silicon Mach-Zehnder interferometers, Opt. Express 18 (3) (2010) 1879–1887. Available from: https://doi.org/10.1364/OE.18.001879.
[132] J. Lee, M. Kim, M. Fournier, P. Labeye, C.J. Oton, F. Testa, Athermal and wavelength-trimmable silicon MZI, in: Optical Fiber Communication Conference, OSA Technical Digest (online) (Optica Publishing Group, 2016), paper M3I.7.
[133] Q. Deng, L. Liu, R. Zhang, X. Li, J. Michel, Z. Zhou, Athermal and flat-topped silicon Mach-Zehnder filters, Opt. Express 24 (26) (2016) 29577. Available from: https://doi.org/10.1364/OE.24.029577.
[134] K. Hassan, C. Sciancalepore, J. Harduin, T. Ferrotti, S. Menezo, B.B. Bakir, Toward athermal silicon-on-insulator (de)multiplexers in the O-band, Opt. Lett. 40 (11) (2015) 2641. Available from: https://doi.org/10.1364/OL.40.002641.
[135] P. Xing, J. Viegas, Broadband CMOS-compatible SOI temperature insensitive Mach-Zehnder interferometer, Opt. Express 23 (19) (2015) 24098–24107. Available from: https://doi.org/10.1364/OE.23.024098.
[136] G. Gao, et al., Silicon nitride O-band (de)multiplexers with low thermal sensitivity, Opt. Express 25 (11) (May 2017) 12260. Available from: https://doi.org/10.1364/OE.25.012260.
[137] S. Tao, et al., Athermal 4-channel (de-)multiplexer in silicon nitride fabricated at low temperature, Photonics Res. 6 (7) (2018) 686. Available from: https://doi.org/10.1364/PRJ.6.000686.
[138] Y.-J. Hung, C.-T. Tang, Z.-W. Song, T.-H. Yen, Athermal polysilicon Mach-Zehnder interferometer in bulk CMOS, IEEE Photonics Technol. Lett. 32 (14) (2020) 855–858. Available from: https://doi.org/10.1109/LPT.2020.3000485.
[139] Z. Mohammed, Z. Mohammed, M. Rasras, M. Rasras, Robust broadband athermal 2 × 2 Mach–Zehnder interferometer with sub-wavelength grating adiabatic couplers, Opt. Lett. 46 (15) (2021) 3781–3784. Available from: https://doi.org/10.1364/OL.431300.
[140] L. Chang, Y. Gong, L. Liu, Z. Li, Y. Yu, Low-loss broadband silicon-on-insulator demultiplexers in the O-Band, IEEE Photonics Technol. Lett. 29 (15) (2017) 1237–1240. Available from: https://doi.org/10.1109/LPT.2017.2706765.
[141] H. Xu, L. Liu, Y. Shi, Polarization-insensitive four-channel coarse wavelength-division (de)multiplexer based on Mach–Zehnder interferometers with bent directional couplers and polarization rotators, Opt. Lett. 43 (7) (2018) 1483. Available from: https://doi.org/10.1364/OL.43.001483.
[142] S.-H. Jeong, Y. Tanaka, Silicon-wire optical demultiplexers based on multistage delayed Mach–Zehnder interferometers for higher production yield, Appl. Opt. 57 (22) (2018) 6474. Available from: https://doi.org/10.1364/AO.57.006474.

[143] S.-H. Jeong, D. Shimura, T. Simoyama, T. Horikawa, Y. Tanaka, K. Morito, Si-nanowire-based multistage delayed Mach–Zehnder interferometer optical MUX/DeMUX fabricated by an ArF-immersion lithography process on a 300 mm SOI wafer, Opt. Lett. 39 (13) (2014) 3702. Available from: https://doi.org/10.1364/OL.39.003702.

[144] Liu Ansheng, et al., Wavelength division multiplexing based photonic integrated circuits on silicon-on-insulator platform, IEEE J. Sel. Top. Quantum Electron. 16 (1) (2010) 23–32. Available from: https://doi.org/10.1109/JSTQE.2009.2033454.

[145] F. Horst, W.M.J. Green, S. Assefa, S.M. Shank, Y.A. Vlasov, B.J. Offrein, Cascaded Mach-Zehnder wavelength filters in silicon photonics for low loss and flat pass-band WDM (de-)multiplexing, Opt. Express 21 (10) (May 2013) 11652–11658. Available from: https://doi.org/10.1364/OE.21.011652.

[146] D. Munk, et al., Eight-channel silicon-photonic wavelength division multiplexer with 17 GHz spacing, IEEE J. Sel. Top. Quantum Electron. 25 (5) (2019) 1–10. Available from: https://doi.org/10.1109/JSTQE.2019.2904437.

[147] L. Han, B.P.-P. Kuo, A. Pejkic, N. Alic, S. Radic, 50GHz silicon cascaded Mach-Zehnder wavelength filter and automatic phase error correction, in: Optical Fiber Communication Conference (OFC) 2019, San Diego, CA, 2019, p. W3B.3. https://doi.org/10.1364/OFC.2019.W3B.3.

[148] Q. Xu, B. Schmidt, J. Shakya, M. Lipson, Cascaded silicon micro-ring modulators for WDM optical interconnection, Opt. Express 14 (20) (2006) 9431–9436. Available from: https://doi.org/10.1364/OE.14.009431.

[149] D.T.H. Tan, A. Grieco, Y. Fainman, Towards 100 channel dense wavelength division multiplexing with 100GHz spacing on silicon, Opt. Express 22 (9) (May 2014) 10408–10415. Available from: https://doi.org/10.1364/OE.22.010408.

[150] Y. London, et al., Energy efficiency analysis of comb source carrier-injection ring-based silicon photonic link, IEEE J. Sel. Top. Quantum Electron. 26 (2) (2020) 1–13. Available from: https://doi.org/10.1109/JSTQE.2019.2934121.

[151] H. Shen, et al., Eight-channel reconfigurable microring filters with tunable frequency, extinction ratio and bandwidth, Opt. Express 18 (17) (2010) 18067–18076. Available from: https://doi.org/10.1364/OE.18.018067.

[152] X. Zheng, et al., A tunable 1x4 silicon CMOS photonic wavelength multiplexer/demultiplexer for dense optical interconnects, Opt. Express 18 (5) (2010) 5151–5160. Available from: https://doi.org/10.1364/OE.18.005151.

[153] C.-H. Chen, et al., A WDM silicon photonic transmitter based on carrier-injection microring modulators, in: 2014 Optical Interconnects Conference, May 2014, pp. 121–122. https://doi.org/10.1109/OIC.2014.6886109.

[154] S. Manipatruni, L. Chen, M. Lipson, Ultra high bandwidth WDM using silicon microring modulators, Opt. Express 18 (16) (2010) 16858. Available from: https://doi.org/10.1364/OE.18.016858.

[155] M. Pantouvaki, et al., "8 × 14Gb/s Si ring WDM modulator array with integrated tungsten heaters and Ge monitor photodetectors, in: Optical Fiber Communication Conference (2014), Paper Th1C.5, March 2014, p. Th1C.5. https://doi.org/10.1364/OFC.2014.Th1C.5.

[156] M. Wade, et al., TeraPHY: a chiplet technology for low-power, high-bandwidth in-package optical I/O, IEEE Micro 40 (2) (2020) 63–71. Available from: https://doi.org/10.1109/MM.2020.2976067.

[157] M. Hammood, L. Chrostowski, N.A.F. Jaeger, Silicon Photonic Bragg Grating Devices, CRC Press, 2021, pp. 223–247. Available from: https://doi.org/10.1201/9780429292033-13.

[158] C. Elachi, C. Yeh, Frequency selective coupler for integrated optics systems, Opt. Commun. 7 (3) (1973) 201−204. Available from: https://doi.org/10.1016/0030-4018(73)90010-2.
[159] K. Ikeda, M. Nezhad, Y. Fainman, Wavelength selective coupler with vertical gratings on silicon chip, Appl. Phys. Lett. 92 (20) (May 2008) 201111. Available from: https://doi.org/10.1063/1.2936862.
[160] A.E. Afifi, et al., Contra-directional pump reject filters integrated with a micro-ring resonator photon-pair source in silicon, Opt. Express 29 (16) (2021) 25173−25188. Available from: https://doi.org/10.1364/OE.431921.
[161] M. Hammood, A. Mistry, H. Yun, M. Ma, L. Chrostowski, N.A.F. Jaeger, Four-channel, silicon photonic, wavelength multiplexer-demultiplexer with high channel isolations, in: 2020 Optical Fiber Communications Conference and Exhibition (OFC), March 2020, pp. 1−3.
[162] F.O. Afzal, et al., Monolithically fabricated subwavelength grating filters for O-band MUX/DEMUX applications, in: 2020 Conference on Lasers and Electro-Optics (CLEO), May 2020, pp. 1−2.
[163] W. Shi, et al., Ultra-compact, flat-top demultiplexer using anti-reflection contra-directional couplers for CWDM networks on silicon, Opt. Express 21 (6) (2013) 6733−6738. Available from: https://doi.org/10.1364/OE.21.006733.
[164] D. Dai, J.E. Bowers, Silicon-based on-chip multiplexing technologies and devices for Peta-bit optical interconnects, Nanophotonics 3 (4−5) (2014) 283−311. Available from: https://doi.org/10.1515/nanoph-2013-0021.
[165] D. Liu, H. Xu, Y. Tan, Y. Shi, D. Dai, Silicon photonic filters, Microw. Opt. Technol. Lett. 63 (9) (2021) 2252−2268. Available from: https://doi.org/10.1002/mop.32509.
[166] D. Dai, L. Liu, L. Wosinski, S. He, Design and fabrication of ultra-small overlapped AWG demultiplexer based on α-Si nanowire waveguides, Electron. Lett. 42 (7) (2006) 400−402.
[167] T. Fukazawa, F. Ohno, T. Baba, Very compact arrayed-waveguide-grating demultiplexer using Si photonic wire waveguides, Jpn. J. Appl. Phys. 43 (5B) (2004) L673. Available from: https://doi.org/10.1143/JJAP.43.L673.
[168] W. Bogaerts, et al., Compact wavelength-selective functions in silicon-on-insulator photonic wires, IEEE J. Sel. Top. Quantum Electron. 12 (6) (2006) 1394−1401. Available from: https://doi.org/10.1109/JSTQE.2006.884088.
[169] Q. Fang, et al., WDM multi-channel silicon photonic receiver with 320 Gbps data transmission capability, Opt. Express 18 (5) (2010) 5106−5113. Available from: https://doi.org/10.1364/OE.18.005106.
[170] W. Bogaerts, et al., Silicon-on-insulator spectral filters fabricated with CMOS technology, IEEE J. Sel. Top. Quantum Electron. 16 (1) (2010) 33−44. Available from: https://doi.org/10.1109/JSTQE.2009.2039680.
[171] S. Cheung, T. Su, K. Okamoto, S.J.B. Yoo, Ultra-compact silicon photonic 512 × 512 25 GHz arrayed waveguide grating router, IEEE J. Sel. Top. Quantum Electron. 20 (4) (2014) 310−316. Available from: https://doi.org/10.1109/JSTQE.2013.2295879.
[172] P. Cheben, et al., A high-resolution silicon-on-insulator arrayed waveguide grating microspectrometer with sub-micrometer aperture waveguides, Opt. Express 15 (5) (2007) 2299−2306. Available from: https://doi.org/10.1364/OE.15.002299.
[173] Y. Zhang, et al., Foundry-enabled scalable all-to-all optical interconnects using silicon nitride arrayed waveguide router interposers and silicon photonic transceivers, IEEE J. Sel. Top. Quantum Electron. 25 (5) (2019) 1−9. Available from: https://doi.org/10.1109/JSTQE.2019.2910415.

[174] Y. Hida, et al., 400-channel 25-GHz spacing arrayed-waveguide grating covering a full range of C- and L-bands, in: OFC 2001. Optical Fiber Communication Conference and Exhibit. Technical Digest Postconference Edition (IEEE Cat. 01CH37171), March 2001, vol. 3, p. WB2. https://doi.org/10.1109/OFC.2001.928320.

[175] S. Kamei, M. Ishii, A. Kaneko, T. Shibata, M. Itoh, N × N cyclic-frequency router with improved performance based on arrayed-waveguide grating, J. Light. Technol. 27 (18) (2009) 4097−4104. Available from: https://doi.org/10.1109/JLT.2009.2022415.

[176] D. Melati, et al., Compact and low crosstalk echelle grating demultiplexer on silicon-on-insulator technology, Electronics 8 (6) (2019). Available from: https://doi.org/10.3390/electronics8060687.

[177] P.J. Bock, et al., Demultiplexer with blazed waveguide sidewall grating and sub-wavelength grating structure, Opt. Express 16 (22) (2008) 17616−17625. Available from: https://doi.org/10.1364/OE.16.017616.

[178] P.J. Bock, et al., Demonstration of a curved sidewall grating demultiplexer on silicon, Opt. Express 20 (18) (2012) 19882−19892. Available from: https://doi.org/10.1364/OE.20.019882.

[179] J. Brouckaert, W. Bogaerts, P. Dumon, D.V. Thourhout, R. Baets, Planar concave grating demultiplexer fabricated on a nanophotonic silicon-on-insulator platform, J. Light. Technol. 25 (5) (2007) 1269−1275.

[180] F. Horst, W.M.J. Green, B.J. Offrein, Y. Vlasov, Echelle grating WDM (de-)multiplexers in SOI technology, based on a design with two stigmatic points, Silicon Photonics Photonic Integr. Circuits, vol. 6996, May 2008, pp. 195−202. Available from: http://doi.org/10.1117/12.781232.

[181] D. Feng, et al., Fabrication insensitive echelle grating in silicon-on-insulator platform, IEEE Photonics Technol. Lett. 23 (5) (2011) 284−286. Available from: https://doi.org/10.1109/LPT.2010.2102347.

[182] Q. Wilmart, et al., A versatile silicon-silicon nitride photonics platform for enhanced functionalities and applications, Appl. Sci. 9 (2) (2019). Available from: https://doi.org/10.3390/app9020255.

[183] D. Dai, X. Fu, Y. Shi, S. He, Experimental demonstration of an ultracompact Si-nanowire-based reflective arrayed-waveguide grating (de)multiplexer with photonic crystal reflectors, Opt. Lett. 35 (15) (2010) 2594−2596. Available from: https://doi.org/10.1364/OL.35.002594.

[184] S. Chen, Y. Shi, S. He, D. Dai, Compact monolithically-integrated hybrid (de)multiplexer based on silicon-on-insulator nanowires for PDM-WDM systems, Opt. Express 23 (10) (May 2015) 12840−12849. Available from: https://doi.org/10.1364/OE.23.012840.

[185] Y.A. Yilmaz, A.M. Alpkilic, A. Yeltik, H. Kurt, Inverse design of efficient and compact 1 × N wavelength demultiplexer, Opt. Commun. 454 (2020) 124522. Available from: https://doi.org/10.1016/j.optcom.2019.124522.

[186] T. Horikawa, et al., Extraction of SOI thickness deviation based on resonant wavelength analysis for silicon photonics devices, in: 2017 IEEE SOI-3D-Subthreshold Microelectronics Technology Unified Conference (S3S), October 2017, pp. 1−3. https://doi.org/10.1109/S3S.2017.8308740.

[187] Y. Xing, J. Dong, S. Dwivedi, U. Khan, W. Bogaerts, Accurate extraction of fabricated geometry using optical measurement, Photonics Res. 6 (11) (2018) 1008−1020. Available from: https://doi.org/10.1364/PRJ.6.001008.

[188] D. Zhu, et al., Integrated photonics on thin-film lithium niobate, Adv. Opt. Photonics 13 (2) (2021) 242−352. Available from: https://doi.org/10.1364/AOP.411024.

[189] H. Xu, et al., Electro-optic activity in excess of 1000 pm V − 1 achieved via theory-guided organic chromophore design, Adv. Mater. 33 (45) (2021) 2104174. Available from: https://doi.org/10.1002/adma.202104174.
[190] F. Eltes, et al., An integrated optical modulator operating at cryogenic temperatures, Nat. Mater. 19 (11) (2020). Available from: https://doi.org/10.1038/s41563-020-0725-5.
[191] T.-J. Lu, et al., Aluminum nitride integrated photonics platform for the ultraviolet to visible spectrum, Opt. Express 26 (9) (2018). Available from: https://doi.org/10.1364/OE.26.011147.
[192] S. Han, T.J. Seok, N. Quack, B.-W. Yoo, M.C. Wu, Large-scale silicon photonic switches with movable directional couplers, Optica 2 (4) (2015) 370. Available from: https://doi.org/10.1364/optica.2.000370.

This page intentionally left blank

Coherent interconnects for data centers

6

Clint Schow[1], *Mark Filer*[2] *and Christopher Doerr*[3]
[1]Department of Electrical and Computer Engineering, University of California Santa Barbara, Santa Barbara, CA, United States, [2]Omniva LLC, Cloud Technology, Santa Clara, CA, United States, [3]Aloe Semiconductor, Inc., Middletown, NJ, United States

6.1 Data center interconnect: system requirements

The term DCI, an acronym for data center interconnect, is generally accepted to refer to the optical interconnects between physical data centers that traverse distances greater than 2 km [1,2]. This chapter will focus on DCI consistent with this definition, considering links spanning 2−100 km. In this space, the links can utilize either intensity-modulated direct detection (IMDD), or coherent detection, with examples of successful product offerings of both varieties. Dense wavelength-division multiplexing (DWDM) in the C-band, centered around 1550 nm, is utilized for the 10- to 100-km links commonly referred to as access or edge DCI. There is also an emerging application for lower cost and power links for 2- to 10-km distances, or campus interconnects, that would likely use gray optics, a term for links that do not utilize WDM and therefore don't require the precise specification of operating wavelength. Within datacenters, for link distances below 2 km, the O-band centered around 1310 nm is generally favored. Fig. 6.1 presents the segmentation of DCI according to link distance [2].

Over the past several years, there has been a shift in how cloud-scale providers (CSPs) conceive and build their datacenters [1]. In the earlier days of the cloud, operators built huge monolithic Mega datacenters to service a local geographical area.

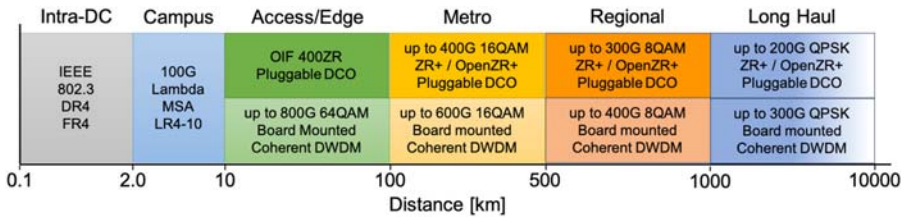

Figure 6.1 Segmentation of datacenter links by distance (*PSM* = parallel single mode; *CWDM4* = course wavelength-division multiplexing with four wavelengths; *PAM4* = pulse amplitude modulation 4-level; *QPSK* = quadrature phase shift keying; *8QAM* = 8-state quadrature amplitude modulation; *16QAM* = 16-state quadrature amplitude modulation). *Source*: Adapted from R. Nagarajan, I. Lyubomirsky, O. Agazzi, Low power DSP-based transceivers for data center optical fiber communications, J. Light. Technol. 2021 [3].

Integrated Photonics for Data Communication Applications. DOI: https://doi.org/10.1016/B978-0-323-91224-2.00005-9
© 2023 Elsevier Inc. All rights reserved.

Over the last several years CSPs have tended to shift toward regional architectures, consisting of smaller, highly interconnected installations within a geographical area to facilitate rapid deployment and efficient scaling. Latency requirements, on the order of a few milliseconds or less, limit the longest interconnect distances in regional architectures to 60−100 km, dictated by time-of-flight through the fiber. Examples of regional architectures for Google, Amazon Web Services, Microsoft, and Alibaba are shown in Fig. 6.2.

As an example, the distribution of link distances for Microsoft's DCI is shown in Fig. 6.3A, where the longest links are 80 km [1]. Distributing datacenters within a region also provides geographic diversity and resiliency to a single site failure.

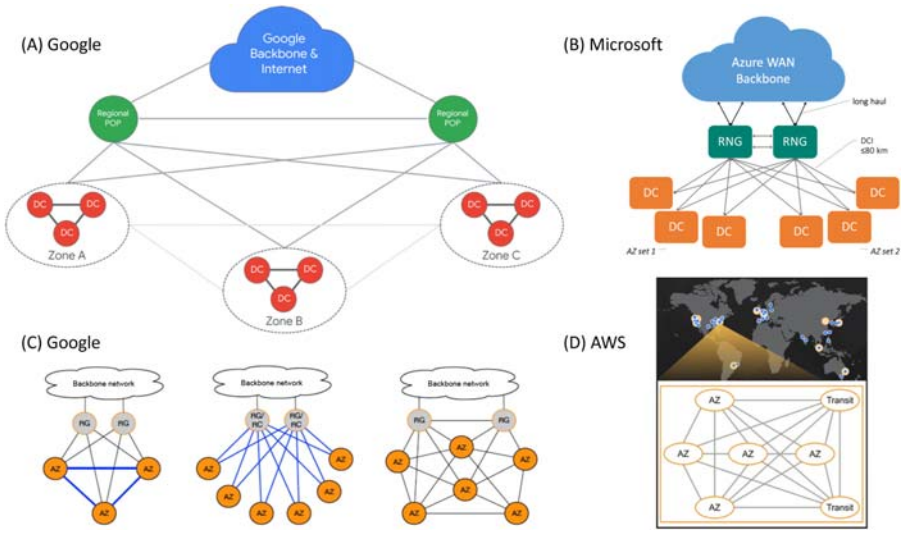

Figure 6.2 Regional datacenter architectures as implemented by (A) Google, (B) Microsoft, (C) Alibaba, and (D) Amazon Web Services (*AZ* = application zone; *RNG* = regional network gateway; *WAN* = wide area network; *DC* = datacenter; *POP* = point of presence).

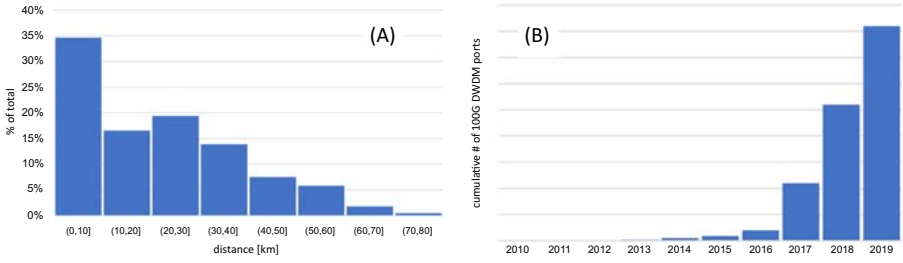

Figure 6.3 (A) Distribution of DCI span lengths and (B) number of deployed 100-G DWDM ports as a function of time for Microsoft's DCI networks.

The transport used for DCI within a region is DWDM due to the large amount of traffic to be supported, but there are typically enough available fiber connections to relegate spectral efficiency to a secondary concern. The move to regional architectures made DCI much more critical, with an explosion in the number of deployed DCI links starting in 2017, as illustrated in Fig. 6.3B.

6.2 Direct detection and coherent links for DCI

In today's optical links, there are two main types of optical signal generation and detection. The first, and oldest, is IMDD, as shown in the upper left portion of Fig. 6.4. In this format, all the information is encoded in the intensity level of the light. Examples are on−off keying (OOK) and pulse amplitude modulation (PAM), which is simply a version of OOK with intermediate levels. The signal is generated by intensity modulation and is detected by a photodiode (PD). The second type of generation and detection is coherent, as shown in the lower and right portions of

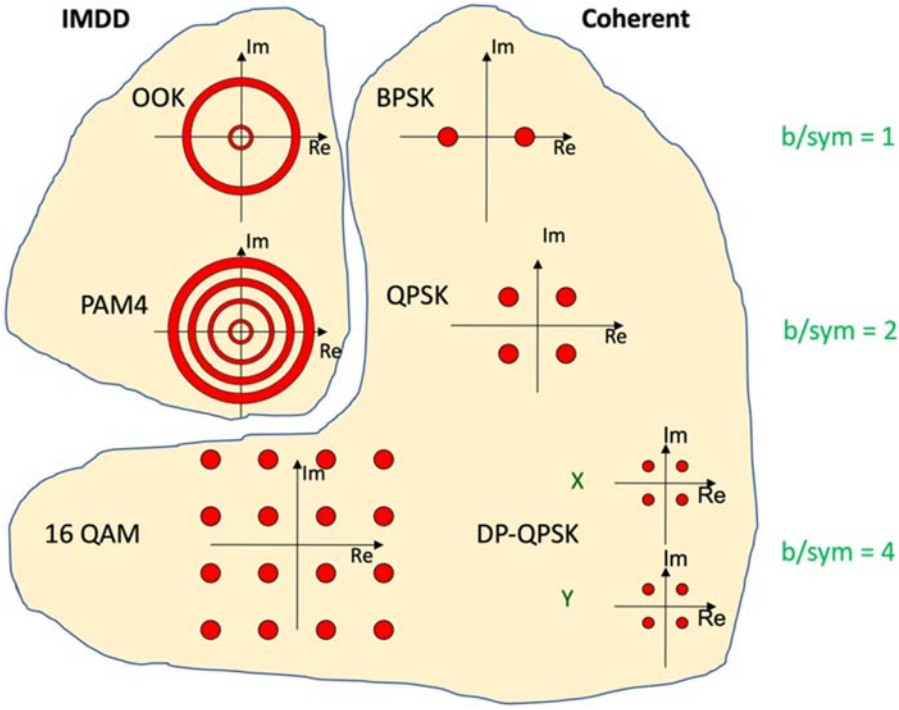

Figure 6.4 Intensity-modulated direct detection and coherent constellations (*OOK* = on−off keying, *PAM4* = pulse amplitude modulation 4-level; *BPSK* = binary phase shift keying; *QPSK* = quadrature phase shift keying; *DP-QPSK* = dual polarization QPSK; *8QAM* = 8-state quadrature amplitude modulation; *16QAM* = 16-state quadrature amplitude modulation).

Fig. 6.4. In this format, information can be encoded not only in the intensity but also the phase and polarization. The signal is generated by phase and amplitude modulators, usually together called a vector modulator, and there is a set for each polarization. The signal is detected by interfering the signal with continuous-wave (CW) light, called a local oscillator, and different relative phases of interference, for each polarization, are sent to separate PDs. A digital signal processor (DSP) then calculates the transmitted signal based on this information.

Fig. 6.5 shows analytical sensitivity calculations for PAM4 and coherent, both running at ~50 Gbaud, based on the equations below. P_b is the bit−error rate, R is the net PD responsivity (including fiber coupling loss, etc.), and N_0 is the total noise power that includes transimpedance amplifier (TIA) thermal noise, shot noise, ADC quantization, and nonlinearity. At 10^{-6} bit−error ratio (BER), in this example coherent exhibits ~12 dB better sensitivity and also delivers 4 × the bit/s. However, coherent tends to exhibit a higher error floor, mainly because it must perform significant MIMO demultiplexing (polarization and I and Q), which is more sensitive to bit quantization in the analog-to-digital converters and any nonlinearity in the system. Note also that coherent has a 2 × shallower BER versus P_{avg} slope than IMDD, which is because coherent and IMDD detect the optical field and power, respectively.

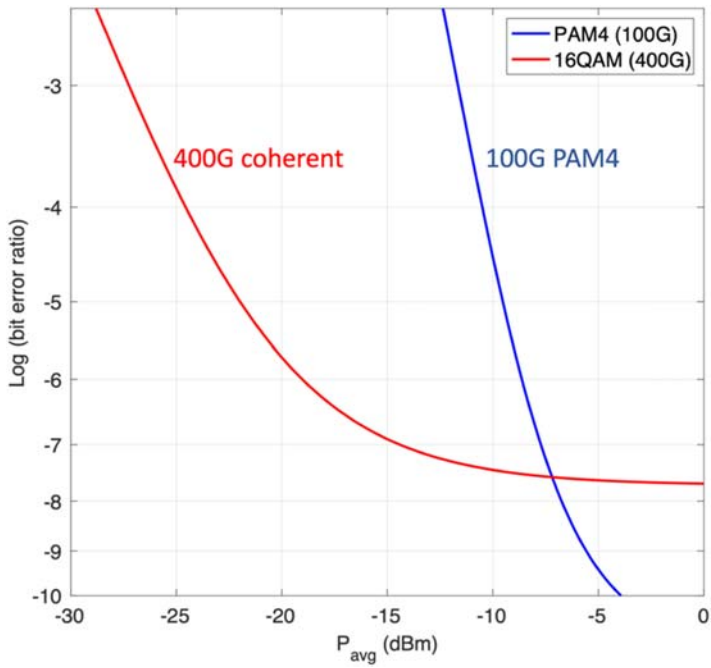

Figure 6.5 Calculated bit−error ratio versus received average power of 400-G coherent (16-state quadrature amplitude modulation) and 100-G pulse amplitude modulation 4-level. Both formats have a symbol rate of ~50 Gbaud.

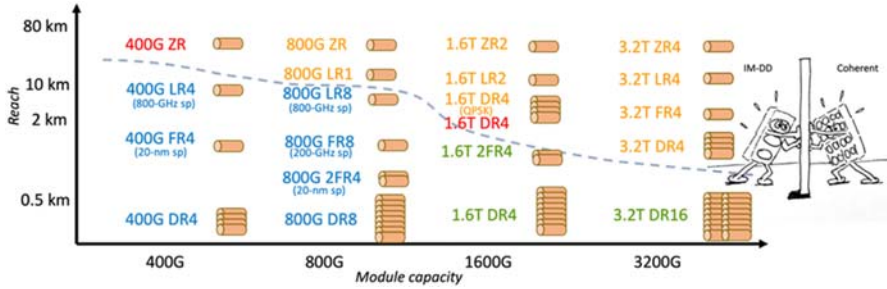

Figure 6.6 Link distance versus rate for various intensity-modulated direct detection and coherent modules. The red and orange colors indicate 50- and 100-Gbaud coherent, respectively, and the blue and green colors indicate 50- and 100-Gbaud PAM4, respectively. The cylinders indicate the number of fibers per communication direction. Coherent assumes 16-state quadrature amplitude modulation unless otherwise noted. *PAM4*, pulse amplitude modulation 4-level.

The higher floor and shallower slope favor using a stronger forward error correction (FEC) for coherent compared to IMDD.

$$P_b = \frac{3}{8}\mathrm{erfc}\left(\frac{P_{\mathrm{OMA}}R}{6\sqrt{(N_0)}}\right)$$

$$P_b = \frac{3}{8}\mathrm{erfc}\left(\sqrt{\frac{P_{\mathrm{sigavg}}P_{\mathrm{LO}}R_{\mathrm{sig}}R_{\mathrm{LO}}}{40N_0}}\right)$$

Because coherent has a larger number of bits per symbol and can compensate for fiber dispersion electronically, it is used primarily for longer links, today greater than ~40 km, and IMDD (specifically PAM4) is used for shorter links, as shown in the module distance vs capacity diagram in Fig. 6.6. Also, coherent is used in the C-band, where fiber loss is low, while IMDD is used in the O-band, where fiber chromatic dispersion is low. Because of coherent's higher complexity, it has larger size, cost, and power consumption than IMDD and today cannot compete at short reaches. As the module rate increases, coherent is expected to move to shorter distances because the penalty from chromatic dispersion grows for IMDD, and the larger number of bits per symbol of coherent means fewer fibers and/or wavelength channels.

6.2.1 Coherent optical devices

Coherent optics requires a great deal of complexity, which is ideally suited for optical integration. Figs. 6.7 and 6.8 show a comparison between IMDD and coherent transmitters and receivers, respectively. The coherent component count is significantly higher, and skews and losses must all be well balanced.

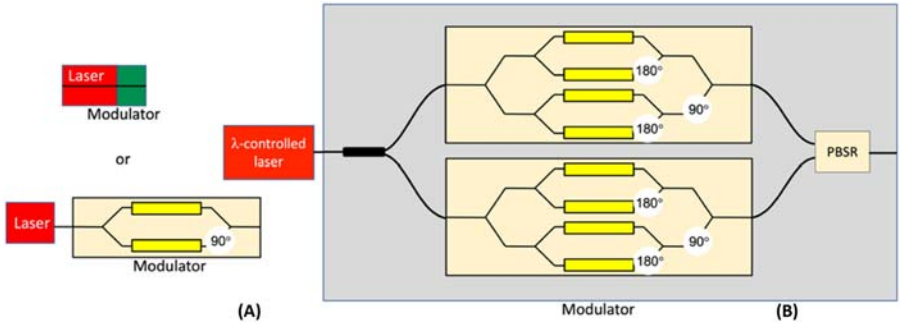

Figure 6.7 Block diagrams of (A) intensity-modulated direct detection and (B) coherent transmitter optical front ends. The intensity-modulated direct detection transmitters can be constructed as an electro-absorption modulated laser (top), or a discrete laser coupled into a Mach−Zehnder modulator (bottom). Coherent transmitters typically employ a laser with precise wavelength control coupled into two nested Mach−Zehnder modulators that generate in-phase (*I*) and quadrature phase (*Q*) signal components for two separate channels carrying distinct data that are then multiplexed onto two orthogonal polarizations for transmission through fiber. The PBSR performs the polarization multiplexing operation. *PBSR*, polarization beam splitter rotator.

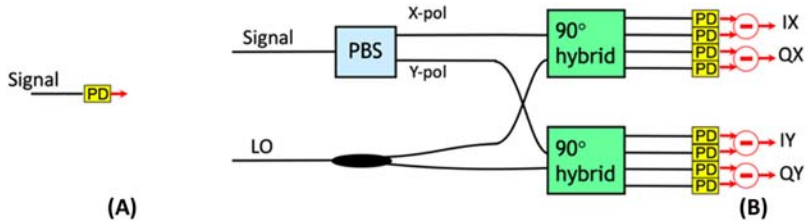

Figure 6.8 Block diagrams of (A) intensity-modulated direct detection and (B) coherent receiver optical front ends. An intensity-modulated direct detection receiver only requires a single high-speed photodiode to convert the incoming optical signal into a corresponding electrical current. Coherent optical front ends employ a PBS to separate the two channels (*X* and *Y*) of the polarization multiplexed incoming signal. The *X* and *Y* channels are then separately mixed with the local oscillator to recover the in-phase (*I*) and quadrature (*Q*) components of the signals for detection by high-speed photodiodes. *PBS*, polarization beam splitter rotator.

6.2.2 Evolution of DCI link technologies from intensity-modulated direct detection to coherent

For DCI, the critical metrics are cost, energy efficiency, density, operational simplicity, and speed of deployment. Traditional long-haul coherent technologies are typically overkill for DCI applications as they offer unneeded features and performance, primarily related to overcoming fiber impairments not found in DCI, all of which consume power, increase footprint, and add to cost. For the 100-G generation, such

considerations led to the development of DWDM PAM4 link technology that fit within the power (4.5 W) and size envelope of the pluggable QSFP-28 form factor. These "ColorZ" modules, produced by Inphi (now Marvell), utilize two wavelengths, each modulated with two bits per symbol (PAM4) at 28 Gbaud to yield a per-lambda data rate of 100 Gb/s, and provide a reach up to 120 km at an aggregate bandwidth of 4Tb/s per fiber pair. Fig. 6.9 shows a photograph of a ColorZ module, typical PAM-4 eye diagram, and the optical spectrum of the module transmitter output and after 120-km transmission and EDFA amplification.

A comparison of the optical spectrum and constellations for the ColorZ versus conventional dual-polarization quadrature phase shift keying 100-G coherent technology is presented in Fig. 6.10. Some notable tradeoffs are apparent. To reach

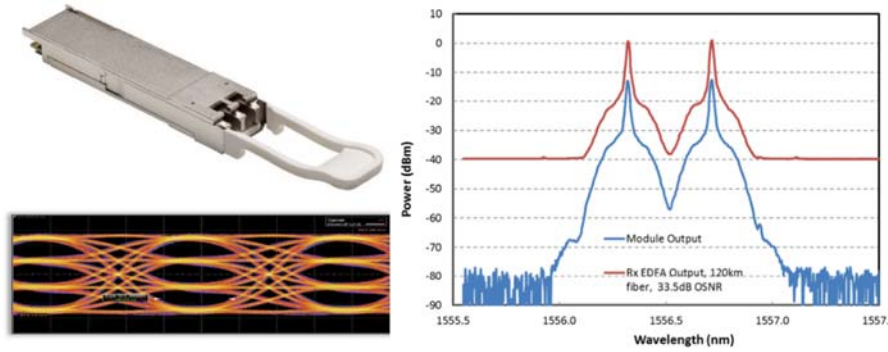

Figure 6.9 ColorZ module, PAM-4 eye diagram, and transmitter output optical spectrum before and after transmission and amplification.

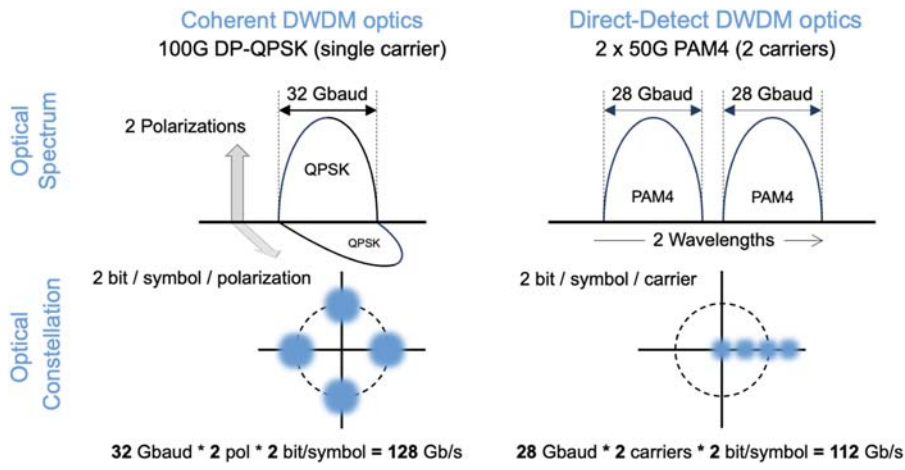

Figure 6.10 Comparison between conventional dual-polarization quadrature phase shift keying 100-G coherent technology and 100-G ColorZ intensity-modulated direct detection.

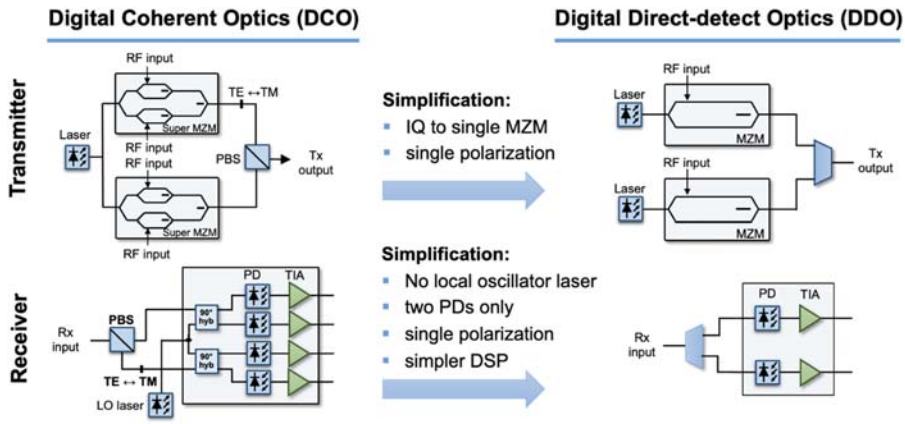

Figure 6.11 Block diagram view of required photonic devices comparing digital coherent optics to digital direct-detect optics.

100 G, the coherent implementation required polarization multiplexing and two bits per symbol spanning both amplitude and phase. The IMDD solution was simpler in terms of only using amplitude and a single polarization, but two carriers (lasers) were needed to reach the full 100-G data rate. The spectral footprint of the ColorZ solution was also a factor of $\sim 2\times$ broader, but this tradeoff is not as critical in DCI due to the relative abundance of fiber.

The relative simplicity of the ColorZ modules compared to the 100-G generation coherent offerings designed for long-haul transport, resulted in significant advantages in cost, density, and power. The significant reduction in complexity for ColorZ at the photonic device level can be seen in the comparison presented in Fig. 6.11, with substantial simplifications for ColorZ at both the transmitter and receiver sides. However, a resulting tradeoff from the reduced DSP capability in the PAM4 modules was the need for optical dispersion compensation, as opposed to the DSP-based electrical compensation found in coherent interconnects. The ColorZ modules leveraged the capabilities of Si photonics to implement a highly integrated photonic integrated circuit (PIC) that includes Mach−Zehnder modulators (MZMs), wavelength multiplexers, and demultiplexers, and a polarization-diverse receiver with high-speed Ge detectors.

As the industry moves from 100 to 400 G and beyond, the scaling advantages of coherent signaling have become clear in comparison to multilevel IMDD. At 400 G, in contrast to 100 G, the added complexity of coherent transmitters and receivers become a favorable tradeoff to meet bandwidth demands. Scaling to 16-QAM modulation enables 4 bits per symbol and polarization multiplexing doubles bandwidth, leading to a per-wavelength bandwidth of $8\times$ the baud rate, enabling a one-laser, 60-Gbd architecture to deliver 400 Gbps. Fig. 6.12 illustrates the DP-16QAM signaling employed to reach 400 Gbps and shows a projection of power consumption that illustrates near parity at 400 G, followed by a crossover point where coherent DWDM offers better efficiency compared to multilevel IMDD.

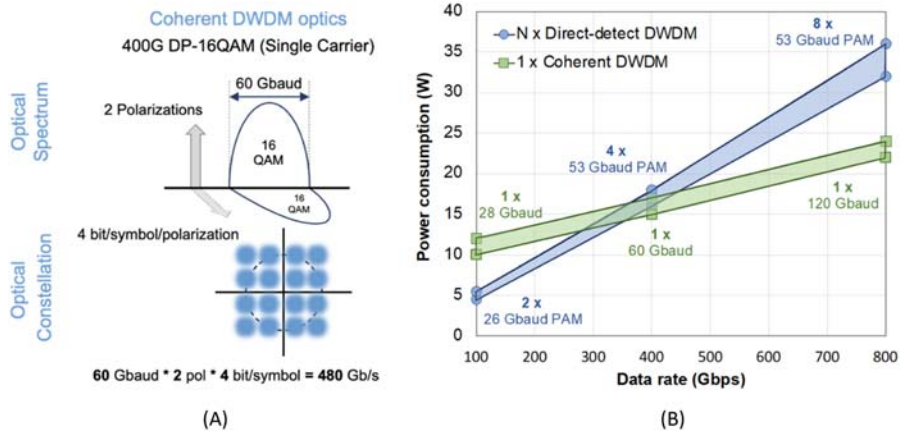

Figure 6.12 (A) Illustration of DP-16QAM signaling used in 400ZR and (B) power consumption projection of coherent dense wavelength-division multiplexing versus intensity-modulated direct detection (pulse amplitude modulation 4-level) links as a function of data rate.

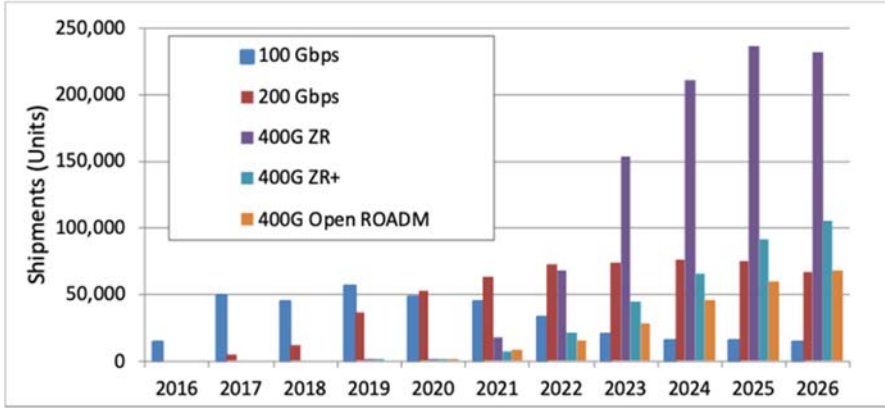

Figure 6.13 Market projections illustrating the rise of 400ZR.
Source: Courtesy of Lightcounting.

For DWDM IMDD, the parallel 1550-nm lasers quickly begin to dominate the module power envelope when comparing to a single carrier, 800-G coherent equivalent. Upper and lower bounds in the plot are representative uncertainty in assumptions on component power consumption.

To enable a 400-G coherent ecosystem, the optical internetworking forum (OIF) undertook the 400ZR standard, which was published in March 2020 [4]. Modules built to the 400ZR standard, which is the first DCI-optimized coherent link technology, are expected to become dominant in the marketplace, displacing IMDD technologies that were viable in the 100-G generation. Fig. 6.13 illustrates this projection,

with 400ZR volumes rapidly growing to significantly exceed legacy 100- and 200-Gbps shipments.

Currently, 400ZR products consume ~16 W while providing 400-Gbps connections over 120 km [3]. A key piece of meeting the power requirements for 400ZR is reducing the power dissipation of DSP by only supporting the capabilities needed for purpose-built DCI applications. In general, the power efficiency of coherent modules has consistently and dramatically improved in the decade since initial deployment at a rate of roughly 56% per year, due in large part to continued CMOS technology scaling [5]. The energy efficiency of 400ZR is 40 pJ/bit, a factor of >30× better than initial coherent solutions introduced a decade ago. An initial module offering employs a low-power DSP implemented in 7-nm CMOS, but the DSP function still accounts for roughly half of the total power consumption of the module—a testament to the importance of CMOS to coherent link scaling.

6.3 Photonic-integrated circuits for coherent DCI

6.3.1 Photonic-integrated circuit material platforms

The three main material platforms used in coherent transmitters today are $LiNbO_3$, the Group III−V materials, and the Group IV materials. Today's coherent receivers are implemented in the Group III−V and Group IV materials. Some also utilize a combination of materials, possibly along with glass waveguides (planar lightwave circuits) or microoptics. These materials have different advantages and disadvantages, and the choice often depends on the application.

6.3.2 Coherent transceivers

Coherent transmitters require demanding specifications. First, they are very complex, requiring phase and amplitude balancing of four MZMs. They require high linearity, otherwise error floors for more advanced formats, such as 64-QAM, become untenable. Finally, they require a large peak-to-peak phase swing to mitigate their large optical loss.

The large optical loss comes from the following: first, $I-Q$ combining adds 3-dB loss to any coherent transmitter. Second, to achieve full transmission through a MZM for the outermost constellation points, coherent requires 2π peak-to-peak differential phase, as compared to half that for IMDD. Third, coherent has its bias point at zero transmission, as compared to IMDD which has its bias point at 3-dB transmission, as shown in Fig. 6.14. The average power transmissions through the modulator for IMDD (PAM4) and coherent (16-QAM), respectively, as a function of peak-to-peak phase are given in the equations below (the coherent equation includes the $I-Q$ combining loss). As one can see, the modulator loss is independent of swing for IMDD but depends approximately quadratically on the peak-to-peak swing for coherent. ϕ_{pp} is the value at Nyquist (half baud rate), not low frequency. A typical $\phi_{pp} = 1.5$ rad at

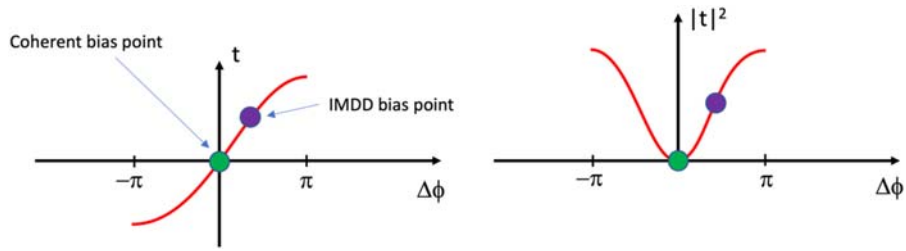

Figure 6.14 Plots of transmission through a Mach−Zehnder modulator versus differential phase. Left shows amplitude transmission, and right shows power transmission.

Figure 6.15 The first $I-Q$ modulator, which was in GaAs.

Nyquist, which gives $T = -14$ dB. This loss is very significant.

$$T = \frac{1}{2}$$

$$T = \frac{5}{18}\sin^2\left(\frac{\phi_{pp}}{4}\right)$$

The very first coherent modulator was in GaAs [6], as shown in Fig. 6.15. This modulator was intended for differential-quadrature-phase-shift keying (DQPSK) and was only a single polarization, but it could have been used for coherent, as well. It used the Pockel's effect, which is a pure crystalline electro-optic effect. The Pockel's effect is very weak, and thus early modulators tended to be $\sim$10 cm long. Later $LiNbO_3$ modulators were deployed, also using the Pockel's effect [7]. Interestingly, today there is renewed interest in Pockel's effect modulators using $LiNbO_3$, but this time as a thin film surrounded by oxide on a silicon substrate. With thin-film $LiNbO_3$, the optical field is more tightly confined; thus, the electrodes can be closer together. This increases the electric field for a given voltage allowing for a shorter modulator. For example, the modulators in the coherent transmitter in Fig. 6.16 have a length of only 1.3 cm [8].

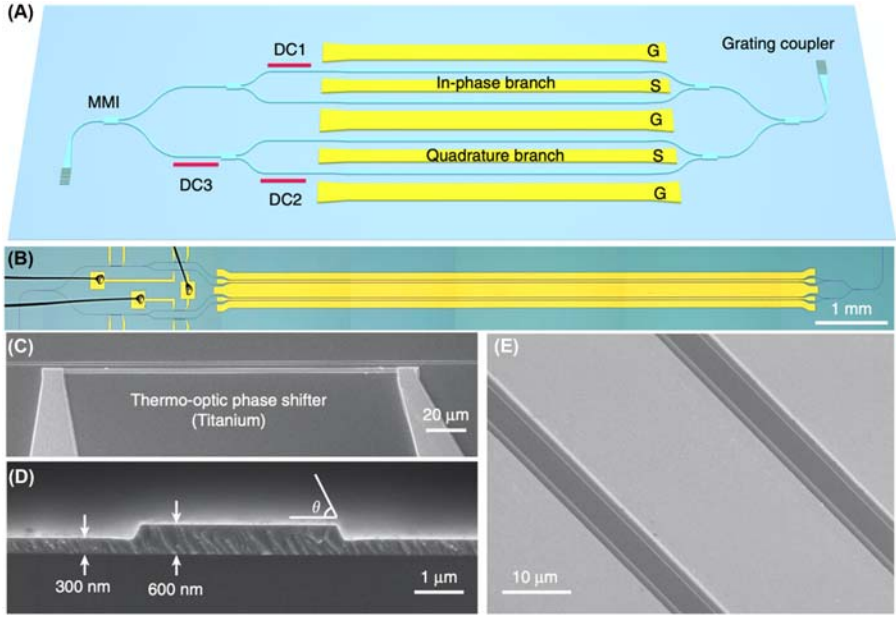

Figure 6.16 Thin-film $LiNbO_3$ coherent modulator.

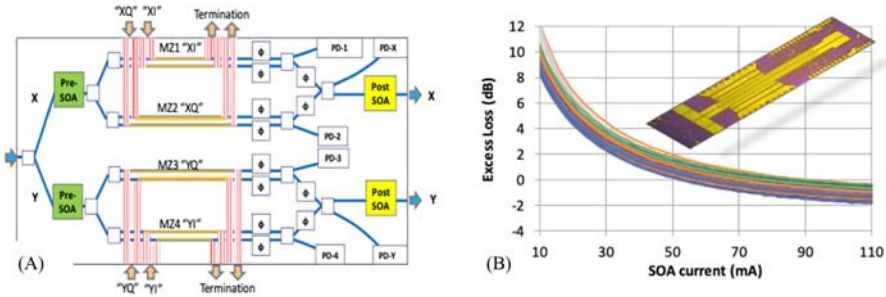

Figure 6.17 InP coherent modulator with integrated optical gain.

InP modulators typically rely on the quantum-confined Stark effect (QCSE). This effect is significantly stronger than the Pockel's effect, resulting in coherent modulators around 2 mm long or even shorter. The main drawback to the QCSE is that it is wavelength and temperature dependent. Thus, it requires both temperature and bias control with wavelength. Also, InP is a relatively expensive material and is only available in small wafers. InP modulators have the additional advantages of having a high bandwidth, with traveling-wave modulators demonstrating bandwidths in excess of 70 GHz, and the ability to monolithically integrate with optical gain. An early coherent modulator in InP is given in Ref. [9] and one with integrated optical gain is shown in Fig. 6.17 [10].

Silicon modulators rely on the carrier depletion effect [11,12]. This effect is strong, with the added benefit of being temperature and wavelength insensitive, but unfortunately it contributes relatively high optical loss. The first silicon coherent modulator was described in [13].

Coherent receivers require multiple optical components integrated together with precise matching of magnitude and phase, including a polarization splitter, hybrids, monitor PDs, and high-speed PDs. The first single-quadrature coherent receiver was made in InP [14], and the first dual-polarization, dual-quadrature coherent receiver was made in silicon [15]. Significant advantages of silicon include the ability to integrate the polarization splitter and rotator and a high-yield process on large wafers.

A historical progression of coherent modulators and receivers is shown in Fig. 6.18. Modulators and receivers have decreased in size and increased in integration level. Because one laser can be shared between the modulator and receiver in most applications, the modulator and receiver are typically integrated together on a single silicon die [16]. Silicon-photonics versions of tranceivers then progressed from being packaged in a hermetic box, to a nonhermetic box, to ball-grid-array substrates, and most recently are being copackaged with the DSP application specific integrated circuit (ASIC) and the PIC sharing the same substrate [17]. In this case, the laser is kept separate and is fiber-coupled to the modulator/receiver die. Other variations have also progressed, notably in InP in which the laser and the modulator both need to be temperature controlled, and are therefore typically copackaged. Currently, the all-silicon approach is best suited for low-cost and/or high-performance applications, and the InP approach is best suited for applications that demand high transmit optical power in a compact footprint.

Significant differences also exist between the source lasers for IMDD versus coherent. Coherent transceivers require a relatively narrow linewidth laser, typically 300 kHz, depending on the baud rate, format, and dispersion [18]; precise wavelength control, typically $\pm$ 1.5 GHz; and relatively high output power, typically 15 dBm. Also, because DCI usually uses DWDM, coherent transceivers require a tunable laser.

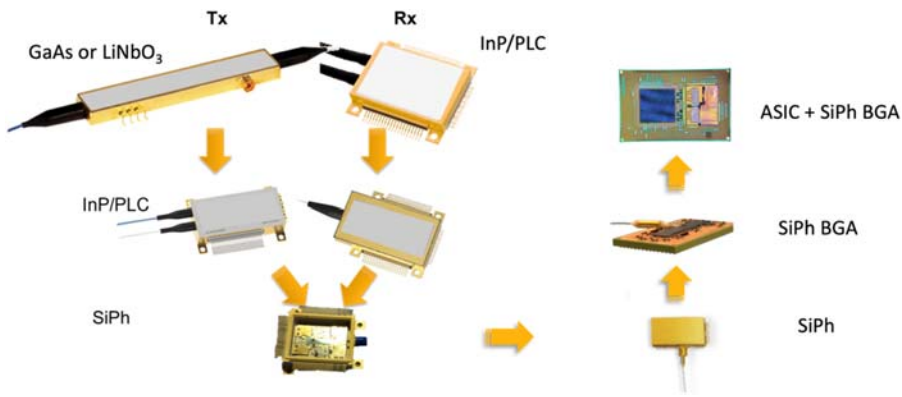

Figure 6.18 Progression of integrated coherent optics.

Except for output power, the requirements for coherent sources are markedly different from IMDD lasers which have no linewidth requirement and only require coarse wavelength control, typically ± 6.5 nm. The output power requirement is about the same.

Because of the demanding laser requirements for coherent, the highest cost component in typical coherent transceivers is the laser. Work is underway to reduce cost, such as by using silicon photonics as an external cavity for the laser [19].

6.3.3 Levels of electrical and optical integration

6.3.3.1 Types of packaging

Fig. 6.19 shows various packaging configurations for photonics and electronics. 2D packaging, also called a multichip module, is placing multiple die on the same substrate. The die may be flip chipped or wire bonded (WB). 2.5D packaging is placing multiple chips on an interposer, usually made of silicon. The advantage of using an interposer is the interconnect density can be significantly higher. 3D packaging is stacking die on top of each other. In 3D packaging, at least one die must have both up and down interconnects. This requires either WB or through-silicon vias.

6.3.3.2 Copackaged optics

There are two main flavors of copackaged optics, or CPO: small ASIC and large ASIC. The small ASIC CPO is copackaging a single ASIC die with a single PIC die. The large ASIC CPO is copackaging one large ASIC die with a plurality of PIC die (Figs. 6.20 and 6.21).

6.3.4 Performance details and limitations

The performance of coherent optics has been proven to improve with integration. One major reason is that coherent requires multiple signals with precisely the same skew, both optical and electrical. Skew error of even a few picoseconds can significantly degrade the performance of >50-Gbaud signals. Optical integration combined with CPO can guarantee subpicosecond skews. Another major reason integration

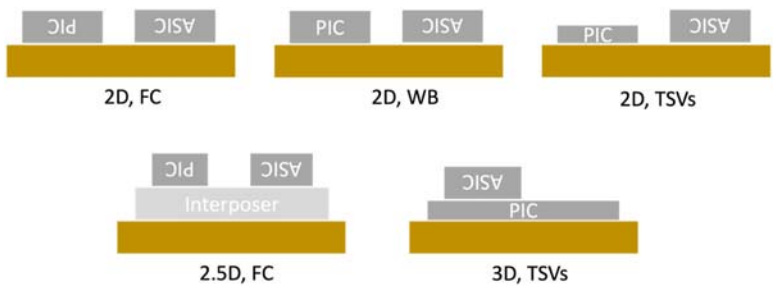

Figure 6.19 Packaging configurations for PICs and ASICs.

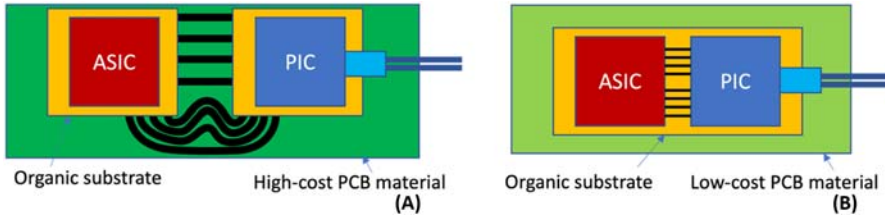

Figure 6.20 Integration to support scaling to 400- to 800-G transceivers. (A) Traditional cohabitation and (B) CPO.

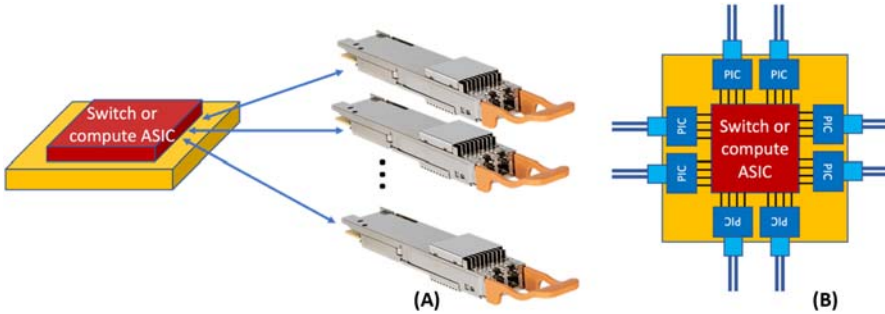

Figure 6.21 Scaling from 25T-100T network switches. (A) Traditional pluggable optics and (B) CPO.

improves performance is that coherent requires precise component matching, both in loss and frequency response. This is because equalization is applied after the *I* and *Q* signals are combined, and if *I* and *Q* have different bandwidths due to component variation, the equalization performance will be limited. Integration guarantees that components next to each other, such as PDs or modulators, have nearly identical performance.

There was initial concern that the slightly nonlinear curve of phase vs voltage in silicon-photonic depletion modulators would penalize silicon photonic coherent transmitters. However, because the modulators are driven push−pull, the nonlinearity is canceled to first order, resulting in negligible performance penalty. A comparison between $LiNbO_3$, which is known to be highly linear, and silicon-photonic modulator performance, presented in Fig. 6.22 demonstrates that the silicon-photonic modulator suffers no significant nonlinear penalty, even with high-order constellations.

6.3.5 Potential for further cost, size, and power reductions for coherent PICs

There are several key future directions for DCI coherent transceivers, the main goal being to reduce the cost, size, and power consumption per transmitted bit/s. An initial

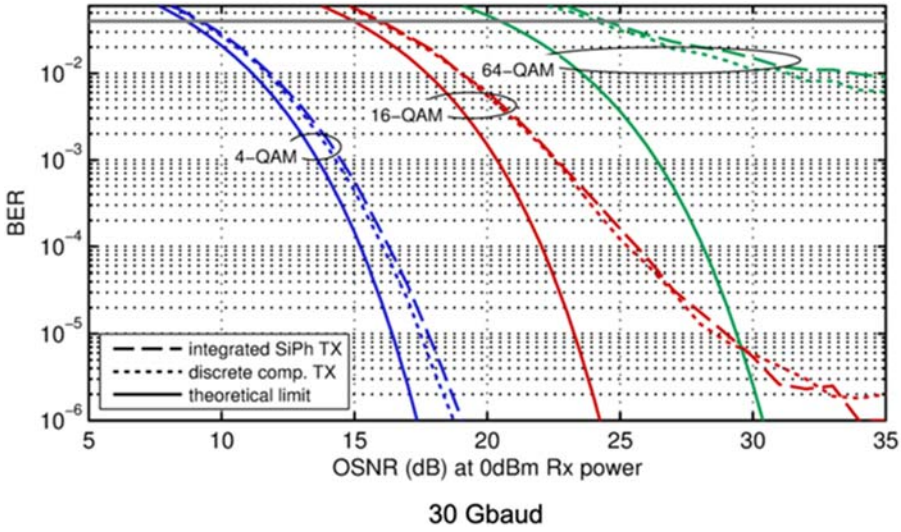

Figure 6.22 Theoretical versus measured performance for $LiNbO_3$ and silicon-photonic transmitters for quadrature phase shift keying, 16-state quadrature amplitude modulation, and 16-state quadrature amplitude modulation.
Source: From Jonas Geyer of Cisco.

thought in coherent was to continue to use higher and higher order constellations. However, a back-of-the-envelope calculation shows this is a limited direction. The equation below shows the conversion from a constellation to bit/s. For example, DP 64-QAM at 63 Gbaud with 27% FEC overhead gives 600 Gb/s. Now consider that you have a 16-QAM system and want to double the bit/s without increasing the baud rate. You would need to go to 256-QAM. Seeing from Fig. 6.22 that the implementation penalty is already very high for 64-QAM, one can see that with today's technology it would be extremely difficult to make 256-QAM work at high baud rate. The main limitations are the ADC effective number of bits and nonlinearity, mostly in the driver and TIA circuits.

$$\frac{\text{bit}}{s} = \text{baud}\frac{\left[\log_2(\#\text{points in } X \text{ pol}) + \log_2(\#\text{points in } Y \text{ pol})\right]}{1 + \text{FEC overhead}}$$

A linear way to increase bit/s is to increase the baud rate. Indeed, this is the main drive in the industry for 800-G ZR, in which the baud rate is doubled from $\sim$60 Gbaud in 400-G ZR to $\sim$120 Gbaud. To achieve a certain baud rate, the transmitter and receiver component bandwidth need to reach the Nyquist frequency, which is half the baud rate. Thus, for a 120-Gbaud system, all the components must have at least 60-GHz bandwidth. Of equal importance to the bandwidth is the number of poles in

the frequency roll-off. For example, it may be acceptable to have a bandwidth slightly less than Nyquist, provided the roll-off is only about two poles.

For the transmitter, the main high-speed components are the DAC, driver, electrical interconnects, and modulator. The DAC and driver should be straightforward to go to ~200 Gbaud. This is because the SNR is high in the transmitter, and thus many techniques, such as peaking, are available. However, these techniques usually result in many poles, which will heavily penalize the performance if Nyquist is not achieved. The electrical interconnects tend to roll off with only 1−2 poles, but it is still important to make them as short as possible as the baud rate increases, especially because of electrical reflections which can cause ripple. Developing a silicon modulator with sufficient bandwidth is a significant challenge in reaching 100 + Gbaud.

The first limitation to the speed of silicon modulators arises from the silicon depletion phase modulation mechanism itself. A typical cross section is shown in Fig. 6.23. In a depletion modulator, electrons and holes must move in and out of the depletion region. The saturated velocity of the holes is lower than that of the electrons, so the speed is limited by that of the holes, which is $\sim 3 \times 10^6$ cm/s. Assuming the depletion width is 0.15 μm, the transit time is 5 ps. This implies ~90-GHz bandwidth, which should allow 200 Gbaud. Other technologies, such as InP's QCSE or $LiNbO_3$'s Pockels effect, do not have such a speed limitation, but do have other drawbacks as previously discussed.

Another potential way to reduce power and increase bandwidth for the transmitter is to use resonant modulators, such as ring modulators [20]. These mainly provide an electrical benefit, increasing electrical impedance at high frequency and thus reducing power consumption and/or voltage for a given current. Ring modulators can be used as either stand-alone amplitude modulators or as phase modulators in MZMs [21]. Also, they can be used in a serial arrangement for wavelength-division multiplexing (WDM), but with a large channel separation due to the typical Lorentzian line shape of resonators. However, they are highly wavelength and temperature sensitive and require control loops to lock them to transmit lasers, and the locking range is limited.

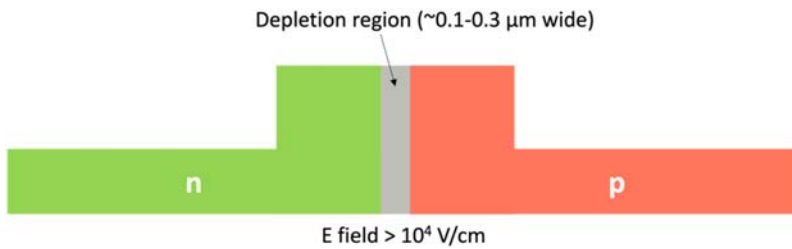

Figure 6.23 Representative cross-section of a high-speed depletion-mode phase shifter used in Si photonic modulators.

For receivers, the main high-speed components are the PDs and TIAs. In Ge PDs, the main speed limitation is typically transit time. This can be continually reduced by reducing the depletion width of the p−n junction. This will increase the capacitance, which can be mitigated by reducing the PD size. The more difficult device to achieve higher bandwidth is the TIA. This is because the TIA is a feed-back device. Today, TIAs are struggling to reach ∼60-GHz bandwidth. Significant transistor performance enhancement or other innovations will be required to reach higher than ∼120 Gbaud.

One method to reduce power and increase bandwidth in the receiver, albeit a small improvement, is to reduce the oversampling in the coherent receiver. In the coherent transmitter, shorter-reach coherent already employs 1 sample per symbol, but in the receiver, usually there is significant oversampling, typically 1.3 samples per symbol at the ADC and internally in the DSP 2.0 sample per symbol. Both the ADC and DSP sampling can be reduced to 1 sample per symbol, at the expense of increased complexity in the timing recovery, performance penalties due to aliasing, and the loss of the ability to electronically compensate for significant optical chromatic dispersion.

Based on the discussion above, one can see that baud rate and constellation point count are both encountering very challenging limitations. Baring their improvement, the main ways left to keep reducing cost, size, and power per bit/s is more integration. This implies more spatial modes and wavelengths per PIC and denser integration with electronics, as shown in Fig. 6.24.

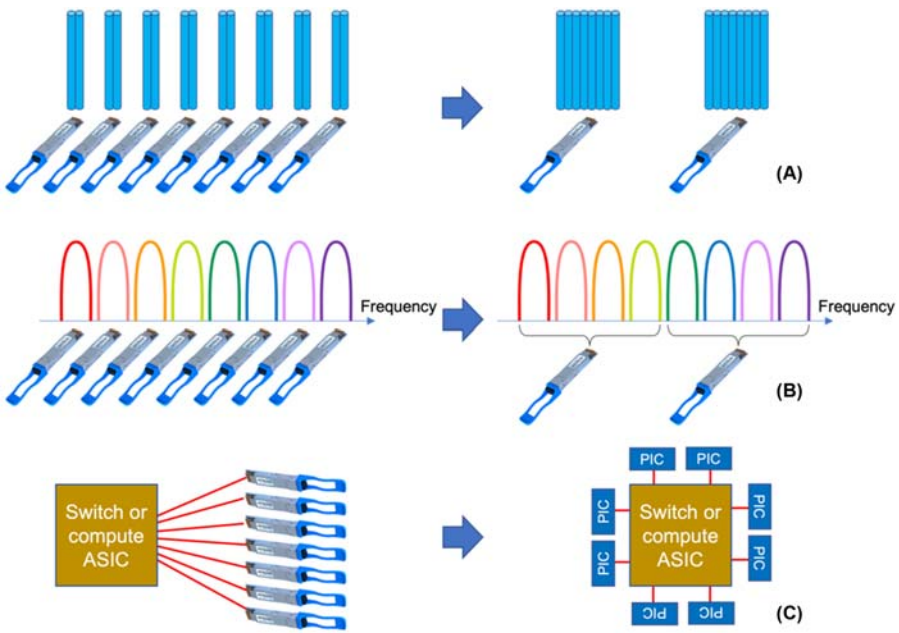

Figure 6.24 Directions for more integration. (A) More spatial modes, (B) more wavelengths, and (C) copackaging with electronic die.

What will ultimately limit such integration is yield. Such a direction may favor silicon photonics over the III−V platform because silicon is a single-atom crystal whereas InGaAsP is a four-atom crystal, requiring precise growth conditions. Other materials, such as thin-film LiNbO3 also required multiatom-type crystals and additionally requires heterogeneous integration techniques such as wafer bonding that require mechanical operations which can introduce particles. For silicon photonics, currently the main yield limitation is the PDs that require Ge epitaxy [22]. However, the industry is making significant progress in this area and yields are already experiencing the multiple 9s required for large-channel-count integration.

6.4 Development paths for volume deployment

6.4.1 The role of standards, consortia, and multi-source agreements

This section presents a case study on CSP participation in standards and consortia to drive development of products tailored to their use cases. CSPs have different (and streamlined) requirements from traditional telcos and enterprises. Examples of this include relaxed specifications: reduced loss and distance requirements for optical modules, or more confined operating temperature ranges due to the tightly controlled data center environment [23]. In addition, products are typically expected to have a shorter life cycle for data center applications, as it's common to perform "forklift upgrades" on rows of servers every 3−5 years for technology refreshes. Likewise, optical modules and transponders for CSP applications need only support Ethernet-framed data and can make do without the added complexities and overhead of optical transport network (OTN) or other legacy protocol support.

The large-scale deployment of 100-G DWDM PAM4 transceiver technology referenced above was viewed as a success in the industry, as those deployments constituted approximately 10% of annual coherent (or coherent equivalent) ports shipped over five consecutive years (see Fig. 6.27) [23]. In addition, such an approach enabled entirely new architectures at scales that were previously unavailable or unaffordable using more traditional means. Recognizing both the success of this targeted product and the large addressable market created by adoption of such architectures, end users and suppliers generally agreed that the next-gen solution should be standardized and interoperable. Particularly for the end users, this serves to create a robust supplier ecosystem and helps ensure that supply chain bottlenecks could not prevent the deployment of regionally distributed data center fabrics.

Two particularly influential CSPs, Google and Microsoft, have been proactive in driving standardized solutions for the 400-G ecosystem. In the OIF, Acacia Communications (now Cisco), Google, Juniper Networks, and Microsoft initiated a project termed 400ZR, to standardize these next-gen optical DCI transceivers.

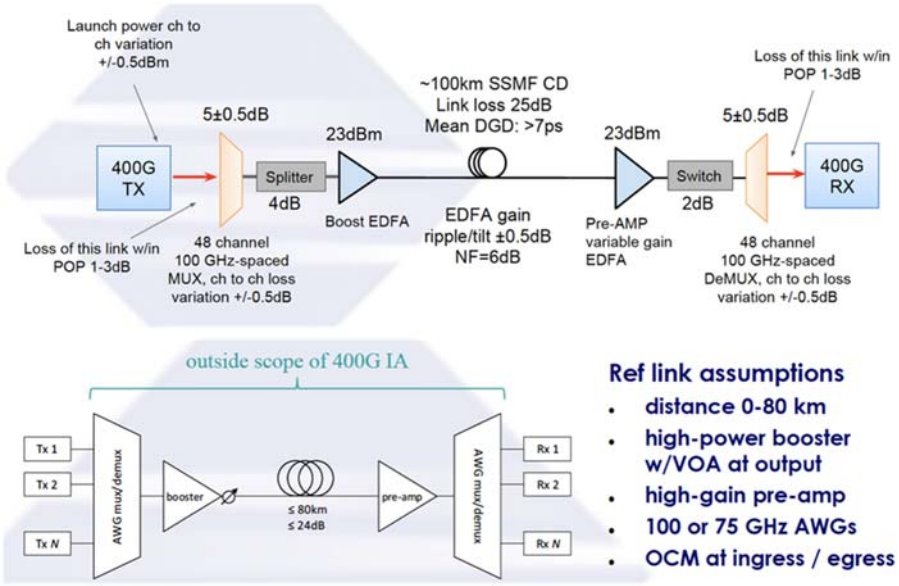

Figure 6.25 Google and Microsoft link requirements shared with OIF in the development of 400ZR.

The concept is similar in principle to that of the DWDM PAM4 technology described above but scaled up to support 400-Gb/s requirements. With 100-G serial PAM4 already pushing the bandwidth limits of the electro-optics, it was quickly determined in the OIF that next-gen solutions for 400 G would need to leverage coherent technologies. However, the drawbacks of ultra-high performance coherent solutions (high power, complexity, and cost) dictated that a different approach be taken with the sub-120 km application space in mind. Google [24] and Microsoft [25] link requirements are shown in Fig. 6.25.

A dedicated, DCI-targeted coherent chipset with minimal features (and therefore minimal power consumption) would need to be developed. Namely, a coherent DSP with limited: (1) dispersion and PMD compensation, (2) baud rate options, (3) flex-modulation support, and (4) FEC modes. Additionally, to make it viable as a switch-pluggable solution like the 100-G DWDM PAM4 modules, it was decided that the optics and electronics (including DSP chip) needed to fit in a power envelope of 15 W and in package sizes to achieve parity with today's solutions ($\geq$32 ports per rack unit (1RU)). Finally, to ensure a robust supplier ecosystem, this technology must all be interoperable, including modulation, framing, pilot tones, and most importantly FEC, which is why the project was carried out in the OIF.

6.4.2 OIF 400ZR specifications

The initial release of the OIF 400ZR implementation agreement (IA) was made publicly available in March 2020. The IA detailed the implementation of a 400-G DCI-focused solution which enabled interoperability across suppliers. High-level features of the 400ZR IA include:

- 400-G Ethernet (400 GbE) only transport supporting 120-km DWDM transmission;
- DP-16QAM modulation format (59.84 Gbaud);
- Standardized concatenated-FEC (CFEC) with 14.8% overhead and 10.8-dB NCG;
- Ethernet framing with simplified FlexO structure;
- 100-GHz nominal channel spacing (75 GHz defined in IA v1.1); and
- 15-W power envelope to support OSFP and QSFP-DD form-factors.

6.4.3 ZR+ derivatives

As the development of the 400ZR IA progressed within the OIF, coherent DSP and module suppliers began mulling the merits of incorporating additional modes to maximize return on investment of a new DSP tape-out and address a larger market. At the same time, interested parties beyond the hyperscalers began showing interest in a lower-cost, multihaul coherent solution, which could fit in pluggable form factors and power envelopes.

The OpenROADM multi-source agreement (MSA) was established in 2015 by AT&T and featured members from other more traditional telcos and enterprises, as well as photonic component, module, and system suppliers [26]. OpenROADM exists to "define interoperability specifications for... transponders and pluggable optics," and the organization quickly targeted the 400ZR effort in OIF to build on for enhanced features sets and performance to address a larger swath of applications and distances [27,28]. In particular, modes were defined to target metro, regional and long-haul applications by implementing variable modulation formats of 16QAM, 8QAM, and QPSK, respectively. Client signal support was extended beyond just 400GE to N × 100GE, 200GE, 400GE, and full OTN support (OTU2/OTU3/OTU4) leveraging FlexO framing. A higher net-coding gain soft-decision FEC was defined with 15.3% overhead, providing an additional coding gain of 0.7 dB compared to CFEC at the expense of 1.4 W more power consumption and increased latency. Furthermore, the symbol rate was increased to 63.1395 Gbaud to support the enhanced performance and OTN client support.

The table that follows summarizes various standardized FEC implementation, showing implementation and performance differences and applications where each have been adopted.

Standard	FEC codes	Data rate	Overhead (%)	Pre-FEC BER	NCG @ 1e − 15 (16QAM) (dB)	Gap-to-Shannon (dB)	Adoption
SC-FEC	Hard-decision staircase FEC	100 G	6.7	4.5e − 3	9.76	1.7	Terastream, 2013 OpenROADM MSA v1.0, 2016 ITU G.709.2 OTU4-SC, 2018 IEEE 802.3-ct, 100GBASE-ZR, 2021
CFEC	Concatenated soft-decision hamming + staircase FEC	400 G	14.8	1.23e − 2	10.78	1.8	OIF 400ZR, 2020 IEEE 802.3cw, TBD ITU G.709.3 FlexO-x-DSH, 2020
OFEC	Soft-decision turbo product code, iterative decoding	200 G 300 G 400 G	15.3	1.94e − 2	11.57	1.1	OpenROADM MSA v3.01, 2019 OpenZR + MSA, 2020 ITU G.709.3 FlexO-x-DO, 2020 OIF 800ZR, TBD

	400ZR	OpenROADM	OpenZR+
Application	Edge DCI	Carrier metro/long haul	Regional DCI
Client support	400 GbE only	100-400 GbE & OTN	100-400 GbE (no OTN)
Reach	≤ 120 km	> 120 km	> 120 km
Form-factor	QSFP-DD/OSFP	CFP2 or other	QSFP-DD/OSFP
SD-FEC	CFEC	OFEC	OFEC
Max power consumption	15 W	Not defined	20 W
Standard	OIF	OpenROADM	Leverages other standards work

Figure 6.26 Comparison of pluggable coherent module standards.
Source: From E. Park, Advancing coherent interoperability with OpenZR + , 4 December 2019. <https://opticalconnectionsnews.com/2019/12/advancing-coherent-interoperability-with-openzr/> [30].

The higher baud rate of the OpenROADM specification, driven mainly the requirement to support OTN clients, pushed the module power envelope beyond what could reliably be packaged in an OSFP or QSFP-DD, necessitating the use of CFP2 form-factor. In response, acknowledging that some hyperscalers or other enterprise users might want the multihaul capabilities of the OpenROADM module but could not tradeoff port density, several module and DSP suppliers formed an MSA called OpenZR + to address these needs [29]. OpenZR + takes the simplified Ethernet-only client support of the 400ZR and the enhanced FEC (oFEC) of the OpenROADM spec to provide a multihaul coherent pluggable, which fits in OSFP or QSFP-DD form-factors, at the expense of higher power consumption (by 2−3 W) when compared to an OIF-compliant 400ZR (Fig. 6.26).

6.4.4 Market impact of IPoDWDM shift

The hyperscaler trend of leveraging switch-pluggable optics using Internet Protocol over DWDM (IPoDWDM) for metro DWDM transport has already had a notable impact on the market, as indicated in Fig. 6.27 [23]. The figure shows the total volume of shipped coherent (or coherent-equivalent) ports for all applications (not just DCI) annually from 2014 to 2020 based on historical data and projecting out to 2025. As indicated, the global deployment of the DWDM PAM4 solution by Microsoft represented approximately 10% of the total market for several years running, further highlighting the massive shift for hyperscalers to move toward more regionally distributed datacenter designs. As 100-G deployments ramp down and 400-G ramps up, the trend toward IPoDWDM is expected to continue, so much so that by 2025, the expected market split between more traditional "embedded" coherent vs pluggable coherent will be approximately 60/40.

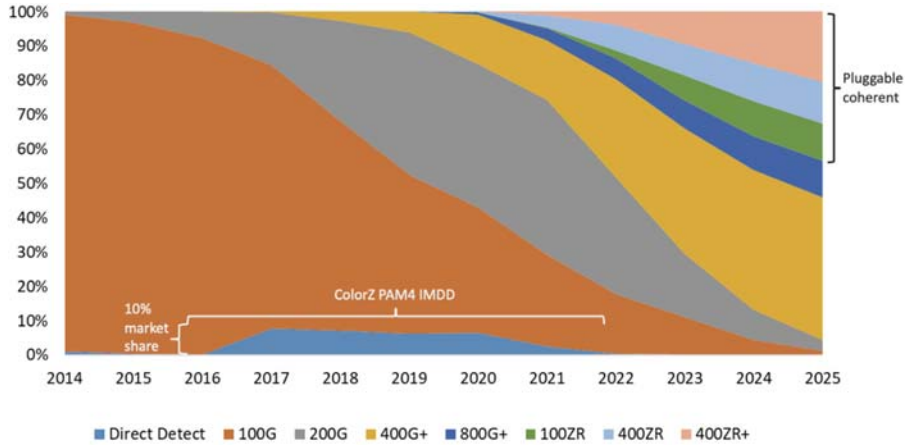

Figure 6.27 Percentage breakdown of the total volume of shipped coherent (or coherent-equivalent) ports [23].

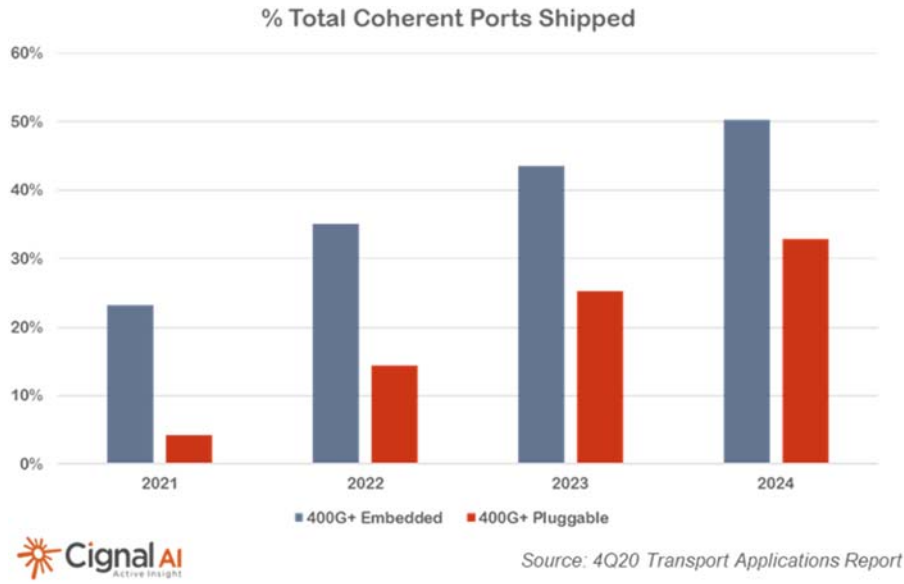

Figure 6.28 Volume breakdown of shipped embedded and pluggable ports.

The chart presented in Fig. 6.28 further highlights the expected trends between embedded and pluggable coherent optics in the 400-G ecosystem [31]. By 2024 the expectation is that the gap between shipped embedded versus pluggable ports will continue to narrow as the number of 400 G ports grows from $<30\%$ of the total volume in 2021 to $>80\%$ in 2024.

6.5 Future directions: outlook beyond 3 years

6.5.1 Interdatacenter interconnects

With the development of the 400ZR standard and looking beyond to the 800-G and 1.6-T generations, coherent link architectures appear to be the path to implementing connections outside the datacenter. The limits imposed by dispersion-induced power fading [32] on IMDD links will become even more severe as baud rates scale beyond 56 GBd, which will continue to favor coherent approaches with their greater tolerance to chromatic dispersion. The use of optical dispersion compensators, which were employed with success in the 100-G generation of PAM-4 IMDD links, will not solve the dispersion issue at higher baud rates. In fact, a 100-km, 56-GBd PAM-4 link would demand chromatic dispersion limits that are lower than the tolerances that could be expected of nominally perfect dispersion compensation [33]. A discussion of chromatic dispersion limits naturally raises the choice of wavelength between the O-band centered around the zero-dispersion wavelength of SSMF at 1310 nm and the C band centered around the minimum loss point of SSMF at 1550 nm. Driven by the extreme bandwidth demands of DCI, the C-band will likely continue to remain the clear choice for 100-km class links due to the overwhelming advantages of optical amplification and support for dense WDM.

In addition to conventional coherent architectures that typically integrate a local oscillator laser within the coherent RX, self-coherent systems have been investigated to reduce complexity and drive down the cost and power of coherent links [34]. Self-coherent approaches transmit an output consisting of the modulated signal combined with a CW tone in an aligned polarization state. The CW component of the transmitter output mixes with the signal at the receiver, thereby performing the LO function, and because the signal and CW propagate through the same fiber path, their polarization remains aligned, potentially removing the need for polarization diverse reception in the receiver.

6.5.2 Intradatacenter interconnects

Within the datacenter, for link distances of less than 2 km, research is underway targeted at reducing the power, complexity, and footprint of coherent transceiver technology to enable it to displace current IMDD solutions as link speeds scale to 400 G and beyond. For optically amplified DWDM DCI, the primary advantages of coherent signaling are high spectral efficiency and tolerance to fiber dispersion, and OSNR tolerance is a critical parameter. Within the datacenter, the O-band around 1310 nm is favored due to minimum fiber dispersion as well as the need for uncooled laser operation that is somewhat easier to achieve compared to 1550 nm. Links inside datacenters are therefore not optically amplified, making the supported optical budget a key metric. Here, coherent links can potentially offer a large improvement in supported link budget due to the signal gain contributed by the local oscillator.

Split-carrier transmitters have been proposed to realize improved power budgets [35]. In this link architecture, the source laser is split such that a fraction of its output undergoes modulation, while the rest of its power bypasses the modulator,

avoiding its high insertion losses. Ideally, this transmitter architecture is implemented on a single PIC, where the modulated signal and unmodulated carrier can be coherently combined in the same polarization state before exiting the PIC. The unmodulated carrier serves as a forwarded LO, boosting the signal when it is mixed with it at the receiver PD. Demonstrations of the split carrier transmitter approach at 112 Gbps/λ, using discrete fiber-coupled components, have yielded improvements of more than 5 dB in link budget [35].

Photonic integration has been leveraged to demonstrate optical phase locked loops (OPLLs) at chip scale [36,37]. This development was a key enabling technology to facilitate "analog coherent" link architectures that can potentially operate without ADCs and DSP, offering significant power savings [38]. In analog coherent links using OPLLs, the LO is locked in frequency and phase to the incoming signal providing direct down-conversion to baseband when using homodyne detection. Employing QPSK modulation in an analog coherent link architecture enables the use of power-efficient analog electronics because after the 90-degree hybrid, the I and Q channels have only two signal levels. This property allows for the use of simple limiting-amplifier based receivers that make hard, binary decisions just like traditional on − off keying nonreturn to zero receivers. Similarly, on the transmitter side, the I and Q channels can be separately driven with two-level NRZ drive electronics. The power efficiency of an OPLL-based analog coherent link can therefore approach that of NRZ IMDD implementations that have been shown to offer superior power efficiency compared to links with multilevel signaling schemes [39].

The larger optical power budgets offered by coherent links also have the benefit of potentially enabling network architectures augmented with photonic routing/switching for performance and efficiency gains. Photonic switches have a long history and continue to be a subject of research around the world. There have been many impressive demonstrations in multiple technology platforms [40−43], but the insertion loss of photonic switches will be a key factor in limiting their adoption for datacenter networks, where typical optical link budgets cannot accommodate even a few extra dB of added loss. The higher sensitivity offered by coherent detection, potentially offering theoretical gains up to 20 dB [44], can be leveraged to allow the integration of passive optical routing (AWGRs) or active photonic switches into datacenter networks. If the potential benefits of such architectures, including per-workload network optimization and disaggregation [45,46], prove to be compelling advantages for datacenter operators, they could further help to drive the adoption of coherent links within datacenters.

6.5.3 What is next in standards?

As of the publication of this manuscript, work is ongoing in the OIF, IEEE and other standards and consortia to define optical PMD definitions and IAs covering next generation data rates. Within the OIF, efforts are underway to define the next evolution of coherent pluggable interfaces, with initial data rates targeting 800 Gb/s [47]. The project has been divided into two separate Implementation Agreements (IAs), one targeting the DWDM application space covered by 400ZR (currently termed 800ZR), and the other to further address campus and intradatacenter gray links for $\leq$ 10-km distances

(termed 800LR). The IEEE formally launched a study group in 2020 called "Beyond 400 Gb/s" to determine which new IEEE project(s) should be launched to support next generation data rates [48]. This led to the creation of the IEEE P802.3df task force which seeks to define interfaces at 800 Gb/s and 1.6 Tb/s (covering derivative physical layer specs for Ethernet frames at 200 and 400 Gb/s) [49].

6.5.4 Summary

The outlook for coherent link technology appears promising, both for inter- and intradatacenter interconnects. A short-term market view for coherent DWDM transceivers is presented in Fig. 6.29, projecting robust growth with growing importance of DCI vs telecom in market percentage. The migration from IMDD to coherent links is now happening in the 400-G generation for DCI, and projections for the uptake of 400ZR in the coming years are strong based on surveys of carriers and hyperscale datacenter operators as shown in Fig. 6.30. Inside the datacenter, the

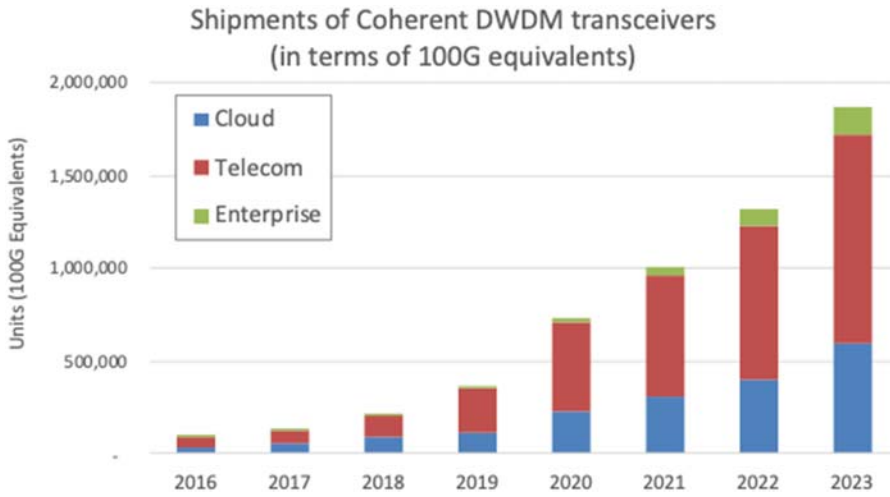

Figure 6.29 Shipments of coherent dense wavelength-division multiplexing transceivers forecasted through 2023.
Source: Courtesy Lightcounting.

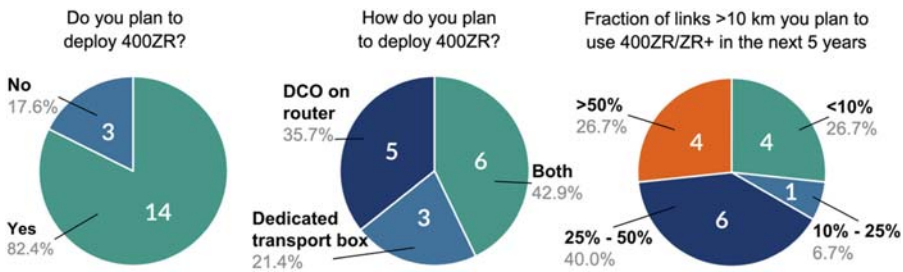

Figure 6.30 Outlook for 400ZR deployment.

migration to coherent may be somewhat further out in time, but scaling IMDD technology to 800 G will be highly, perhaps prohibitively, challenging. At the 1.6- and 3.2-T generations, coherent is projected to offer substantial advantages over IMDD: fewer required lasers with less output power needed per laser, and greater tolerance to dispersion and link penalties [50].

Integration and packaging will continue to be critical in driving down the footprint and power of coherent transceivers. Integration of III − V active components into Si photonic platforms with low coupling loss that can scale to high volumes at low cost will be critical. Similarly, low-loss coupling of many fibers, on the scale of 100, to chip packages will also be required and will need to scale to high volume, high-yield, low-cost packaging. Flip-chip integration will likely be ubiquitous in photonic transceiver engines to minimize their footprint along with packaging parasitics. By definition, technology platforms that monolithically offer electronic and photonic circuits represent the ultimate in electronic/photonic integration [51−53]. Such platforms can support comprehensive built-in testing and high yields through reduced packaging steps, enabling the potential for low cost at high volumes. Of course, continued scaling of CMOS will have a dramatic impact on the energy efficiency of both IMDD and coherent links, with the latter benefiting

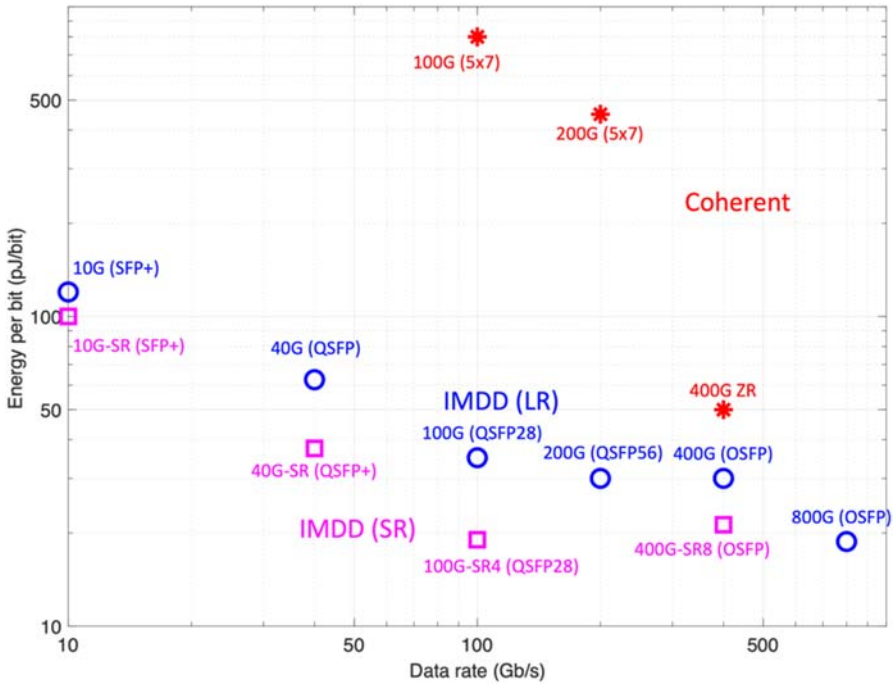

Figure 6.31 Link efficiency projections as data rates move to 800 Gbps.
Source: Data from Cedric Lam of Google X. Zhou, R. Urata, H. Liu, Beyond 1 Tb/s intra-data center interconnect technology: IM-DD OR coherent?, J. Light. Technol. 38 (2) (2020) 475−484 [51].

significantly from DSP scaling in each technology generation. Market forecasting predicts a cross-over from IMDD to coherent driven by bandwidth scaling, and as the crossover approaches, the link energy efficiency advantage for IMDD compared to coherent links is projected to narrow as shown in Fig. 6.31.

There is widespread expectation that integration of optical *I/O* into first-level IC packages will be demanded in the coming years, driven by the unending need for higher-bandwidth network switches. Often referred to as "co-packaged optics [54]," this new paradigm of electronic and photonic integration will require holistic codesign of chip and optical *I/O* and the optimal partitioning of link functional blocks between host ICs and copackaged optical engines. For coherent technologies to be viable for copackaging, improvements in both power efficiency and bandwidth density will be required. Furthermore, the question of where lasers reside for copackaged optics will impact how coherent transceivers are implemented. Currently, there appears to be substantial momentum for external light sources, and common specifications are being developed [55]. Not having access to a "local" LO would certainly impact coherent link architectures and may favor self-coherent approaches that avoid the need for an LO integrated in the coherent receiver.

References

[1] M. Filer, J. Gaudette, Y. Yin, D. Billor, Z. Bakhtiari, J.L. Cox, Low-margin optical networking at cloud scale, OSA J. Opt. Commun. Netw (2019) C94–C107.

[2] R. Nagarajan, M. Filer, F. Yang, M. Kato, T. Rope, J. Stewart, Silicon photonics-based 100 Gbit/s, PAM4, DWDM data center interconnects, IEEE/OSA J Opt. Commun. Netw. 10 (7) (2018) 25–36.

[3] R. Nagarajan, I. Lyubomirsky, O. Agazzi, Low power DSP-based transceivers for data center optical fiber communications, J. Light. Technol (2021).

[4] OIF, OIF-400ZR-01.0 – Implementation Agreement 400ZR, 10 March 2020. <https://www.oiforum.com/wp-content/uploads/OIF-400ZR-01.0_reduced2.pdf>.

[5] F. Frey, R. Elschner, J.K. Fischer, Estimation of trends for coherent DSP ASIC power dissipation for different bitrates and transmission reaches, in: Photonic Networks; 18. ITG-Symposium, Leipzig, Germany, 2017.

[6] R. Griffin, R. Johnstone, R. Walker, S. Wadsworth, A. Carter, M. Wale, Integrated DQPSK transmitter for dispersion-tolerant and dispersion-managed DWDM transmission, in: OFC, 2003.

[7] T. Sakamoto, A. Chiba, T. Kawanishi, 50-Gb/s 16 QAM by a quad-parallel Mach-Zehnder modulator, in: ECOC, 2007.

[8] M. Xu, M. He, H. Zhang, J. Jian, Y. Pan, X. Liu, et al., High-performance coherent optical modulators based on thin-film lithium niobate platform, Nat. Commun. 11 (3911) (2020).

[9] N. Kikuchi, et al., 80-Gb/s low-driving-voltage InP DQPSK modulator with an n-p-i-n structure, IEEE Photonics Technol. Lett. 21 (12) (2009) 787–789.

[10] R.A. Griffin, N.I. Whitbread, S.K. Jones, S.C. Heck, P. Firth, D. Govan, et al., InP coherent optical modulator with integrated amplification for high capacity transmission, in: OFC, 2015.

[11] R.A. Soref, B.R. Bennett, Electrooptical effects in silicon, IEEE J. Quant. Electron. QE-23 (1) (1987) 123–129.

[12] A. Liu, L. Liao, D. Rubin, H. Nguyen, B. Ciftcioglu, Y. Chetrit, et al., High-speed optical modulation based on carrier depletion in a silicon waveguide, Opt. Express 15 (2007) 600−668.
[13] P. Dong, C. Xie, L. Chen, L.L. Buhl, Y.-K. Chen, 112-Gb/s monolithic PDM-QPSK modulator in silicon, Opt. Express 20 (2012) B624−B629.
[14] T.L. Koch, et al., GaInAs/GaInAsP multiple-quantum-well integrated heterodyne receiver, Electron. Lett. 25 (1989) 1621−1623.
[15] C.R. Doerr, P.J. Winzer, Y.-K. Chen, S. Chandrasekhar, M.S. Rasras, L. Chen, et al., Monolithic polarization and phase diversity coherent receiver in silicon, J. Light. Technol. 28 (4) (2010) 520−525.
[16] C. Doerr, L. Chen, D. Vermeulen, T. Nielsen, S. Azemati, S. Stulz, et al., Single-chip silicon photonics 100-Gb/s coherent transceiver, in: OFC, 2014.
[17] C. Doerr, L. Chen, Silicon photonics in optical coherent systems, Proc. IEEE 106 (12) (2018) 2291−2301.
[18] T. Pfau, S. Hoffmann, R. Noe, Hardware-efficient coherent digital receiver concept with feedforward carrier recovery for M-QAM constellations, J. Light. Technol. 27 (8) (2009) 989−999.
[19] X. Huang, C.R. Doerr, C. Qin, J. Heanue, N. Zhu, D. Ton, et al., Silicon-photonic laser emitting tunable dual wavelengths with highly correlated phase noise, Opt. Lett. 46 (2021) 142−145.
[20] Q. Xu, S. Manipatruni, B. Schmidt, J. Shakya, M. Lipson, 12.5 Gbit/s carrier-injection-based silicon micro-ring silicon modulators, Opt. Express 15 (2007) 430−436.
[21] C. Chang, G.d. Valicourt, S. Chandrasekhar, P. Dong, Differential microring modulators for intensity and phase modulation: theory and experiments, J. Light. Technol. 35 (15) (2017) 3116−3124.
[22] T.-Y. Liow, K.-W. Ang, Q. Fang, J.-F. Song, Y.-Z. Xiong, M.-B. Yu, et al., Silicon modulators and germanium photodetectors on SOI: monolithic integration, compatibility, and performance optimization, IEEE J. Sel. Top. Quantam Electron. 16 (1) (2010) 307−315.
[23] Cignal AI, 2Q21 transport customer markets report, 14 September 2021. <https://cignal.ai/2021/09/2q21-transport-customer-markets-report/>.
[24] X. Zhou, T. Hofmeister, V. Dangui, Open coherent target link budgets, OIF contribution oif2017.114.00. <https://oiforum.com>, 2017.
[25] M. Filer, OIF 400G reference links, OIF contribution oif2017.111.00. <https://oiforum.com>, 2017.
[26] A. Fuetsch, Driving the open ROADM: how AT&T is redesigning the optical network, AT&T, 16 March 2016. <https://about.att.com/innovationblog/032216roadm>.
[27] Open ROADM MSA, Open ROADM MSA. <http://openroadm.org/home.html>, 2016.
[28] M.A. Sluyski, OpenROADM_MSA3.01-W-Port-Digital-Specification, 25 June 2019. <https://0201.nccdn.net/1_2/000/000/141/b6c/OpenROADM_MSA3.01-W-Port-Digital-Specification.docx>.
[29] Open ZR + Multi Source Agreement, Open ZR + technical specification, 4 September 2020. <http://www.openzrplus.org/site/assets/files/1075/openzrplus_1p0.pdf>.
[30] E. Park, Advancing coherent interoperability with OpenZR +, 4 December 2019. <https://opticalconnectionsnews.com/2019/12/advancing-coherent-interoperability-with-openzr/>.
[31] Cignal AI, 4Q20 transport applications report, 25 March 2021. <https://cignal.ai/2021/03/4q20-transport-applications-report/>.

[32] U. Gliese, S. Norskov, T.N. Nielsen, Chromatic dispersion in fiber-optic microwave and millimeter-wave links, IEEE Trans. Microw. Theory Tech. 44 (10) (1996) 1716–1724.
[33] X. Chen, C. Antonelli, S. Chandrasekhar, G. Raybon, A. Mecozzi, M. Shtaif, et al., Kramers–Kronig receivers for 100-km datacenter interconnects, J. Light. Technol. 36 (1) (2018) 79–89.
[34] X. Chen, S. Chandrasekhar, P. Winzer, Self-coherent systems for short reach transmission, in: 2018 European Conference on Optical Communication (ECOC), Rome, Italy, 2018.
[35] T. Gerard, Z. Liu, M. Sezer Erkilinc, P. Bayvel, B. Thomsen, D. Lavery, Improved power budget of 112 Gb/s/λ intra-datacentre links using a split-carrier transmitter architecture, IEEE Photonics J. 11 (4) (2019) 1–15.
[36] M. Lu, et al., An integrated 40 Gbit/s optical costas receiver, IEEE J. Light. Technol. 31 (13) (2013) 2244–2253.
[37] P.A. Binetti, et al., Indium phosphide photonic integrated circuits for coherent optical links, IEEE J. Quantum Electron. 48 (2) (2012) 279–291.
[38] J.K. Perin, A. Shastri, J.M. Kahn, Design of low-power DSP-free coherent receivers for data center links, IEEE J. Light. Technol. 35 (21) (2017) 4650–4662.
[39] T. Hirokawa, et al., Analog coherent detection for energy efficient intra-data center links at 200 Gbps per wavelength, IEEE J. Light. Technol. 39 (2) (2021) 520–531.
[40] B.G. Lee, A.V. Rylyakov, W.M.J. Green, S. Assefa, C.W. Baks, R. Rimolo-Donadio, et al., Monolithic silicon integration of scaled photonic switch fabrics, CMOS logic, and device driver circuits, IEEE J. Light. Technol. 32 (4) (2014) 743–751.
[41] K. Tanizawa, K. Suzuki, M. Toyama, M. Ohtsuka, N. Yokoyama, K. Matsumaro, et al., Ultra-compact 32 × 32 strictly-non-blocking Si-wire optical switch with fan-out LGA interposer, Opt. Express 23 (13) (2015) 17599–17606.
[42] T.J. Seok, N. Quack, S. Han, R.S. Muller, M.C. Wu, Large-scale broadband digital silicon photonic switches with vertical adiabatic couplers, Optica 3 (1) (2016) 64–70.
[43] H.R. Mojaver, A. Das, B. Rad, V. Tolstikhin, K.-W. Leong, O. Liboiron-Ladouceur, Scalable SOA-based lossless photonic switch in InP platform, in: 2020 IEEE Photonics Conference (IPC), Virtual Conference, 2020.
[44] G.P. Agrawal, Fiber-Optic Communication Systems, Wiley, New York, 2012.
[45] M. Glick, N.C. Abrams, Q. Cheng, M.Y. Teh, Y.-H. Hung, O. Jimenez, et al., PINE: photonic integrated networked energy efficient datacenters (ENLITENED Program), J. Opt. Commun. Netw. 12 (12) (2020) 443–456.
[46] A.A.M. Saleh, K.E. Schmidtke, R.J. Stone, J.F. Buckwalter, L.A. Coldren, C.L. Schow, The INTREPID program: technology and architecture for next-generation, energy-efficient, hyper-scale data centers [invited], J. Opt. Commun. Netw. (2021).
[47] Optical Internetworking Forum, OIF 800G coherent. <https://www.oiforum.com/technical-work/hot-topics/800g-coherent/>.
[48] IEEE, IEEE launches 802.3™ beyond 400Gb/s ethernet study group to define next high speed threshold for ethernet data rate, 3 December 2020. <https://standards.ieee.org/news/2020/802_3-ethernet-study-group-high-speed-threshold/>.
[49] IEEE, IEEE P802.3df 200 Gb/s, 400 Gb/s, 800 Gb/s, and 1.6 Tb/s ethernet task force, 17 March 2022. <https://www.ieee802.org/3/df/index.html>.
[50] X. Zhou, R. Urata, H. Liu, Beyond 1 Tb/s intra-data center interconnect technology: IM-DD OR coherent? J. Light. Technol. 38 (2) (2020) 475–484.
[51] K. Giewont, et al., 300-mm monolithic silicon photonics foundry technology, IEEE J. Sel. Top. Quantum Electron. 25 (5) (2019) 1–11.

[52] M. Rakowski et al., 45 nm CMOS—silicon photonics monolithic technology (45CLO) for next-generation, low power and high speed optical interconnects, in: 2020 Optical Fiber Communications Conference and Exhibition (OFC), San Diego, CA, USA, 2020.
[53] S. Lischke, D. Knoll, C. Mai, L. Zimmermann, Advanced photonic BiCMOS technology with high-performance Ge photo detectors, in: Proc. SPIE 11088, Optical Sensing, Imaging, and Photon Counting: From X-Rays to THz, San Diego, CA USA, 2019.
[54] Co-Packaged Optics Collaboration. <http://www.copackagedoptics.com/>.
[55] Co-Packaged Optics Collaboration, Co-packaged optics external laser source guidance document. <http://www.copackagedoptics.com/wp-content/uploads/2020/05/CPO-Assembly-Guidance-Doc-V1.0-FINAL.pdf>.

Photonic-integrated circuits for switched network interconnects

7

Saeed Fathololoumi[1], *Yuliya Akulova*[1], *Ashkan Seyedi*[2] *and David Hui*[1]
[1]Intel, Silicon Photonics Product Division, Santa Clara, CA, United States, [2]NVIDIA Corporation, Santa Clara, CA, United States

7.1 Data center interconnect: system requirements

As introduced in Chapter 1, data centers (DCs) are the backbone of digital networks. DCs are warehouse-sized facilities (up to 10 million square feet [1]) with hundreds of rows of racks all connected using optical links forming a large optical network (Fig. 1.3). There are several optical network architectures that are used, including the fat-tree network, which typically has three layers of switches [spine, fabric, and top of rack (TOR)] forming a full mesh topology between the core switches and the racks (Fig. 1.6).

Within a rack, servers or network interface cards (NICs) are typically connected using passive copper cables. For longer connections (up to 100 m), active electrical cables or active optical cables (AOCs) are often used. Active electrical cables do not have any optics, they have electrical repeaters which perform a 3R regeneration of the electrical signal allowing for longer reach in the electrical domain. These active electrical cables have a limited reach of $<$10 m and have even shorter reach at higher data rates. For longer links between racks in a row, AOCs are often used. AOCs convert the electrical data to optical pulses and send them through optical fibers, providing longer reach. An AOC is essentially two short reach (SR) optical transceivers with a multimode (MM) optical cable that is permanently attached between them. These SR transceivers are usually built with vertical cavity surface emitting lasers (VCSEL) technology and discrete receivers, and they have a reach of up to 100 m.

Communication between rows of servers inside the DC typically use pluggable transceivers with detachable optical fibers that are part of the fiber plant. The value of having fixed fiber plants is so that the DC can upgrade the overall bandwidth (BW) without the need to change the fiber infrastructure. Older DCs use MM fibers; hence, they are limited to SR or long reach (LR) multimode transceivers, which use VCSEL technologies. Newer DCs can be much larger with link distances between the TOR and fabric switches longer than 500 m, and links to spine switches are even longer. For these DCs, single-mode fiber (SMF) is often used inside the fiber plant in combination with datacenter reach (DR) and LR optical transceivers. These transceivers are built with either discrete optics, such as directly modulated lasers (DMLs) and externally-modulated lasers (EMLs), or with integrated optics, such as those based on silicon photonic-integrated circuits (SiPICs). SiPIC technologies are gaining popularity because of their smaller size, lower

Integrated Photonics for Data Communication Applications. DOI: https://doi.org/10.1016/B978-0-323-91224-2.00002-3
© 2023 Elsevier Inc. All rights reserved.

power consumption, higher BW, as well as their ability to use the latest CMOS fabrication technology to deliver high yield and low-cost solutions.

This chapter focuses on photonic-integrated circuits (PICs) for switch networks with reaches from 100 m up to 10 km for DC applications. Chapter 8 will focus on PICs to enable optical circuit switching.

7.1.1 Intradata center link requirements and standards

The global DC traffic grows at a compound annual rate of 25% [2], and most of this traffic stays within the DC. The rate of data traffic increase within hyperscale DCs is even higher, close to doubling every two years. The switch and transceiver BWs therefore need to grow at similar rate to keep pace with demand, as shown in Table 1.1. As illustrated in Fig. 7.1, the DC bandwidth of switch node is doubling every two years. Similarly, the scaling of the optical transceivers' transmission capacity is aligned to that of the switch generations and doubles every 2−3 years. The demand for increasing transmission capacity at low-cost, high-energy efficiency, and maximum deployment flexibility resulted in the adoption of pluggable front plate optical transceivers for inter- and intra-DC applications over the past two decades. To overcome fiber transmission impairments at higher modulation rates and speed limitations of optical and electrical components, the optical transceiver architectures rely on scaling along three orthogonal design axes of lane count (fiber and wavelength), lane speed, and bits per symbol.

The IEEE 802.3 Ethernet working group governs the optical and electrical interface standards for each technology generation. It has released specifications covering intensity modulation and direct detection (IMDD) transceiver architectures from 100 to 800 Gbps, scaling lane count from 4 to 8, baud rates from 26 to 53 Gbaud,

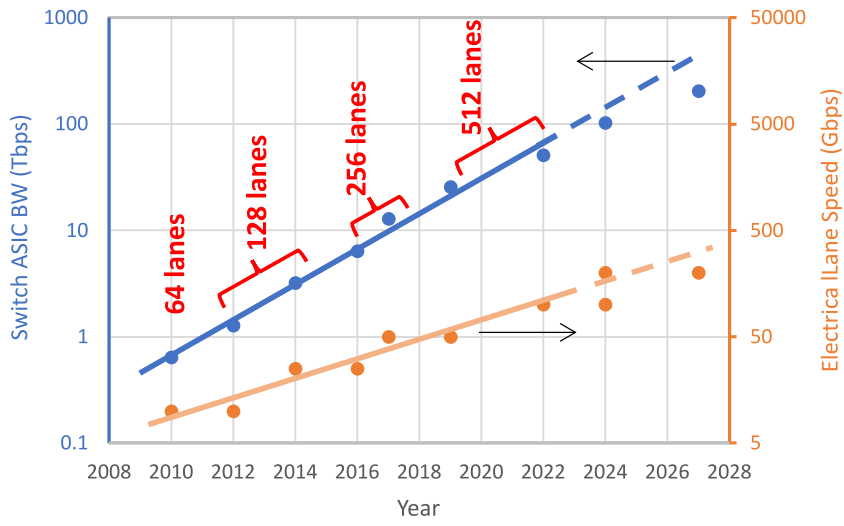

Figure 7.1 Data center bandwidth and transceiver growth [3].

and modulation formats from nonreturn to zero (NRZ) to 4-level pulse amplitude modulation (PAM4) [4]. PAM4 signaling doubles bit per symbol but incurs additional signal-to-noise ratio (SNR) penalty such that low error rate transmission is no longer feasible at high BW and errors must be compensated for using error correction. In typical error compensation schemes, such as forward error correction (FEC), the data are encoded and wrapped in a frame that leads to encoding overhead and added latency of transmission. For 100-G data transmission, the typical KP FEC requires transmission rate increase to 112 G with a latency of ~200 ns.

Table 7.1 summarizes examples of quad small form-factor pluggable (QSFP) transceivers from 40 to 800 G showing how power consumption, power efficiency and bandwidth scale through different generations of these optical transceivers. While 100-G base transceivers have been widely deployed in DCs, 200 and 400 G transceivers are ramping in volume, and 800 G transceivers are already available for early evaluation. However, transceiver power density and complexity are becoming key technical challenges. Currently deployed switches are each 25.6 T, which requires 64 × 400 G transceivers. In current switch systems there are limited space on the faceplate to accommodate all 64 QSFP transceivers. Also, despite energy per bit improving from previous generation, the doubling of data rate means that the power required to operate the switch and transceivers would increase drastically from previous generation. Next generation switch systems at 51.2 Tbps and 102.4 Tbps and their interconnecting optics would face even more power and bandwidth density challenges. Furthermore, starting from the 400-G generation, OSFP pluggable with wider width was introduced to accommodate 8 high-speed *I/O* lanes, further exacerbating the bandwidth density issue.

In switch chassis designed for pluggable modules, the switch ASIC is placed in the middle of the chassis with pluggable transceivers at the faceplate (Fig. 7.2—left panel). This design offers the flexibility to use different types of transceivers (AOC, SR, DR, and LR transceivers) to match the reach requirement of any given port and to allow easy maintenance (replacement of failed units). The electrical interfaces of all these modules need to conform to an electrical specification (e.g., CAUI-x). With the common electrical interface, the optical interface can be changed to accommodate different network requirements. Currently most widely deployed transceivers use LR SERDES to interface with the switch ASIC.

Unfortunately, these switch chassis designs are not power efficient. As data rates increase, electrical losses between the switch and the face plate are also increasing, and transceivers are required to have built-in clock and data recovery (CDR) or SERializer/

Table 7.1 Comparing datacenter requirement for different QSFP generations.

Type	Modulation	Rate (Gbps)	Typ power (W)	Power eff (pJ/b)	BW density (Gbps/cm^2)
40-G QSFP +	NRZ	40	3.5	87.5	2565
QSFP28	NRZ	100	7	70	6411
QSFP56	PAM4	200	7	35	12,823
QSFPDD	PAM4	400	12	30	25,645
QSFPDD 800	PAM4	800	18	22.5	51,290

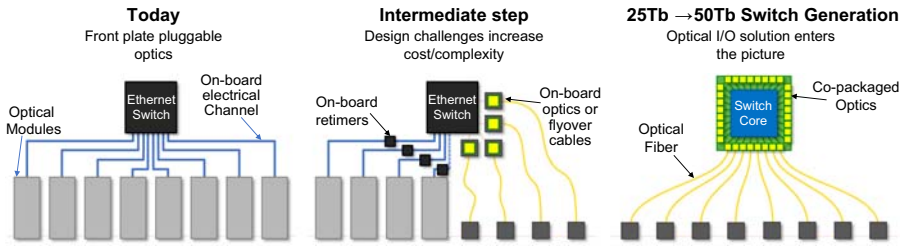

Figure 7.2 Switch chassis design configurations for pluggable optics (left panel), re-timed pluggable optics and on-board optics (middle panel), and copackaged optics (right panel).

DESerializer (SERDES) to clean up the data signals. These CDR/SERDES are power hungry and can consume up to 50% of the total transceiver module power. For longer electrical links, such as those between switch and edge transceivers, retimers are also used to extend the reach, further driving up power consumption. Another issue is that the high-power transceivers are usually located at the front faceplate and warm up the incoming air to the chassis, making it more difficult to cool the switch ASIC inside.

The pace of bandwidth increase will continue, and next generation of pluggable transceivers is expected to deliver 800 Gbps and 1.6 Tb/s transmission capacity with 200 Gbps per lane based on 112-Gbaud PAM4. Such architecture is projected to improve transceiver power efficiency (pJ/bit) and cost ($/bit) since the number of optical channels is reduced by a factor of two relative to 100 Gbps per lane for the same transmission capacity. However, the increased baud rate rapidly reduces achievable optical reach, due to reduction in eye diagram opening, higher noise bandwidth and increased sensitivity to dispersions. It scales down the reach from 10 to $<$3 km for CWDM4 wavelength grid IMDD [5].

With the increase in data rate, electrical channels between switch ASIC and optics face signal integrity (SI) impairments due to bandwidth limitations and increased RF loss. The pluggable systems might require signal retiming to compensate for these impairments. Alternatively, the optics can be placed in the middle of the board to reduce the electrical channel length. Such configuration is called on-board optics (OBO). With the increase of the electrical data rate beyond 100 Gbps, more power is needed to correct for SI impairments even for OBO. Hence the optics should ideally be placed on the same package, as shown on the right panel of Fig. 7.2. This implementation is called copackaged optics (CPO). Currently these new switch architectures are being investigated to reduce power, increase data density, and relieve face plate congestion. CPO modules or engines typically target high BW ($>$1.6 Tbps) and function like many transceivers packaged into one. Although unlike transceivers, engines include only the most essential transceiver elements, such as the PIC and the electrical integrated circuits (EICs). These ICs are assembled using 2.5D or 3D packaging to reduce the required space. All the required voltage regulators are moved onto the host board. The pluggable connector can be replaced by a vertically mounted latch grid array (LGA) connector, and the optical connector is moved from the module to the faceplate of the switch chassis. All these changes make the overall size of the engine small; with the small

size, many engines can be tightly packaged around the perimeter of the switch ASIC and shorten the electrical channels between them from 10s of centimeters to 10s of millimeters, thereby reducing the SERDES driving requirement (from LR and MR SERDES to XSR SERDES) and saving power. Additional chip-level and assembly integration, such as 3D die stacking of PIC and EIC, will further shrink the engines such that they can be placed within millimeters of the switch ASIC. This way the SERDES (which is the highest power consumption component) inside the engine can even be removed, allowing for an unretimed (direct drive) engine [6]. At the faceplate, only the optical connector is needed, hence reducing the faceplate area requirement.

7.1.2 Electrical signal integrity

SI plays an important role in switch architecture. Fig. 7.3 shows a cartoon of a generic transceiver link from host Tx to host Rx. Electrical signals at high-speed suffer ISI penalty from losses in the transmission line as described previously. Currently powerful SERDES are used to compensate for this penalty. One-third to half of the transceiver's power is used by these SERDES. To reduce the power required by these SERDES, the channel needs to be short. Different types of SERDES (shown in Table 7.2) are available, and each is capable of compensating for different ranges of channel length and impairments [7]. If the distance between the switch and the transceiver can be short, lower power SERDES can be used on both sides of the channel. For longer channels (channels routing length longer than the SERDES range), repeaters can be used to extend the reach; thereby limiting the power increase to only a

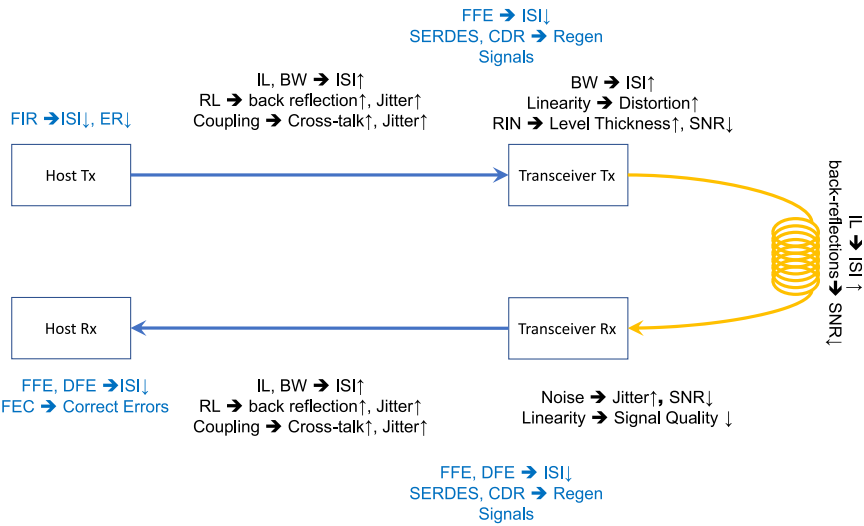

Figure 7.3 Cartoon of generic transceiver link from host Tx to host Rx. FIR, Finite impulse response filter; ISI, inter-symbol interference; ER, extinction ratio; IL, insertion loss; RL, return loss; FFE, feed forward equalization; CDR, clock data recovery; RIB, relative intensity noise; SNR, signal to noise ratio; DFE, decision feedback equalization.

Table 7.2 Different types of SERDES for different ranges of channel length at 56 and 112 G [7].

		14 GHz (OIF-56G)			28 GHz (OIF-112G)			Application
		Reach (mm)	**IL (dB)**	**Power (pJ/bit)**	**Reach*** **(mm)**	**IL (dB)**	**Power (pJ/bit)**	
LR	PAM4	1000	35	6.5	<1000	36	6.5	Backplane
MR	NRZ/PAM4	500	20	3.6	<500	28	4	C2C
VSR	NRZ/PAM4	125	10	2.7	100	21	3	C2M
XSR	NRZ/PAM4	50	4	1.5	50	12	1.5	C2 nearby OE
USR	NRZ	10	1.5	<1	10	2	<1	C2 OE

*Assume Tachyon100/Megtron 7.

few longer channels. Long channel lengths also increase the probability of crosstalk or noise pickup. Some of the reflections in short channels can be mitigated by the FIR at the output of the Tx SERDES. Reflections in longer channel lengths can appear in subsequent bits that are not within the range of the FIR; hence, FIR cannot be used to clean up these channels. Moving the optics closer to the switch (i.e., CPO) allows much shorter channel lengths and thereby increasing the electrical SI and reducing the power consumption and the other penalties.

The effective BW of the link is the overall cascaded BW of all the constituent components. For 100-G NRZ signaling, the Nyquist frequency is 50 GHz. Achieving a cascaded BW of 50 GHz can be very challenging because the overall high-speed link has many components [driver, modulator, transmission line, photodetector (PD), and transimpedance amplifier (TIA)]. To lower the bandwidth requirement for 100 G, the IEEE standards working group changed the modulation format from traditional NRZ to PAM4. This PAM4 modulation format encodes two bits per symbol; therefore, it achieves the same data rate as NRZ but at half the Nyquist bandwidth. But encoding with PAM4 comes with many penalties. By splitting the two-level system into four levels, the signal-to-noise ratio automatically drops by 4.77 dB. Furthermore, to enable the four-level system, the electrical circuitries for driving and amplifying the signals need to have a linear design instead of the limiting design used by NRZ; these circuits are more complex and are more power hungry.

When PAM4 modulation was introduced, different performance metrics were needed to qualify the transmitter for interop with different receivers. First, an ideal 5-tap FFE receiver is chosen as the golden receiver. Then equipment vendors designed the equivalent golden receiver into their scopes for measuring the Tx eye characteristics. The transmitter dispersion eye closure (TDEC) metric from NRZ testing is extended to PAM4 testing as TDEC quaternary (TDECQ) for qualifying the PAM4 Tx Eye. TDECQ is the optical penalty of the Tx eye as compared to an ideal Tx to reach the same bit−error ratio, which is 2E-4 with KP FEC, using the 5-tap FFE golden receiver. Ceq is the noise enhancement factor based on the weights of the 5-tap FFE, and RLM is the linearity metric, which is calculated using the spacings between the four levels of the PAM4 signals.

$$\mathrm{RLM} = \left(3\mathrm{ES}_{1,}3\mathrm{ES}_{2,}2 - 3\mathrm{ES}_{1,}2 - 3\mathrm{ES}_{2}\right),$$

$$\mathrm{ES}_1 = \frac{V_1 - V_{\mathrm{mid}}}{V_0 - V_{\mathrm{mid}}}, \mathrm{ES}_2 = \frac{V_2 - V_{\mathrm{mid}}}{V_3 - V_{\mathrm{mid}}}, V_{\mathrm{mid}} = \frac{V_0 + V_3}{2},$$

Other Tx metrics from NRZ modulation that are also used in PAM4 are outer extinction ratio (ER), outer optical modulation amplitude (OMA), and average optical power (AOP).

7.1.3 Photonic-integrated circuit loss bottlenecks and challenges

The integration of photonic components on a single monolithic chip provides cost and performance advantages for solutions to meet the link requirements mentioned

in the previous section. However, the optical link budget still well exceeds 10 dB. Most of the link budget is usually consumed by three main elements: coupling laser light into PIC, optical modulation, and light coupling between PIC and fiber.

One of the main challenges for PICs (particularly for SiPICs) is to provide enough power across all the desired wavelengths at the input waveguide. The InP PIC benefits from monolithic growth of the gain material within the same platform. However, the multiwavelength products require multiple growth cycles. On the other hand, SiPICs suffer from not having native gain on the same platform. Hence, the light must be generated in an InP based laser and coupled to the Si WG. We will discuss the various methods of laser integration with SiPIC platform in the next section.

Optical modulation is usually achieved in PICs by attenuating the laser light for lower modulation levels. For instance, for NRZ modulation the input power to modulator (excluding the modulator loss) is the "1" level or P1, and the attenuated power level is the "0" level or P0. The ratio of the two powers is defined as the ER (ER = P1/P0), which is in linear scale but can also be express in logarithmic scale as dB. The AOP and OMA are related to the linear ER as follows:

$$\text{AOP} = P_{\text{in}}\left(\frac{\text{ER}+1}{2\text{ER}}\right)$$

$$\text{OMA} = P_{\text{in}}\left(\frac{\text{ER}-1}{\text{ER}}\right)$$

Fig. 7.4 is a plot of AOP and OMA penalties as a function of ER in dB. For example, for an ER of 3 dB, OMA penalty is 3B. On top of this penalty, modulators in the PIC have additional insertion loss, either as a result of the inherent absorption of the modulators or their operating bias points.

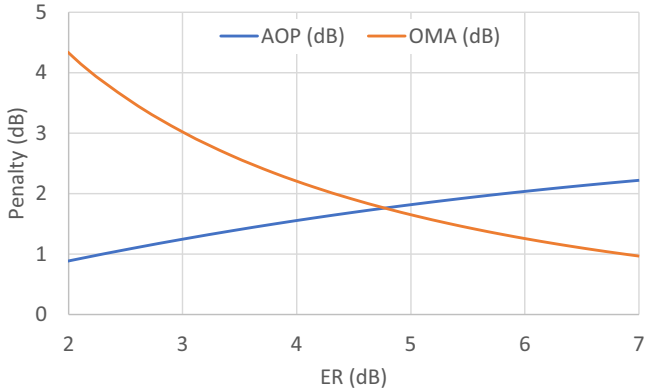

Figure 7.4 Tx average optical power and optical modulation amplitude penalties as a function of modulator extinction ratio.

Traditional transceivers with discrete photonic components or subassemblies are typically designed to couple light between the components and to fibers using free-space optics and/or butt-coupling. Both coupling approaches involve high-precision alignment steps that are sensitive to mechanical drift during assembly and over operating life, which could consume a large portion of the optical link budget. PICs, on the other hand, integrate the required active and passive photonics components on a single chip with self-aligned, interconnecting waveguides. The waveguides can also be designed with mode converters to match the fiber mode for low loss coupling (as discussed in Chapter 5). Also, if lasers are integrated, a PIC is expected to only interface with the fiber once at the output of Tx and once at the input of Rx and benefit from not needing any other coupling interface. Nevertheless, any fiber coupling is expected to mechanically degrade over life, so proper over-life margin must be included in the optical power budget analysis. Recently many foundries have integrated V-grooves to PICs that can provide mechanical support for fibers and hence are expected to suffer lower over-life degradation.

7.2 Integrated photonic building blocks for data center interconnect

PIC technologies for datacenter interconnects are built using either InP or silicon platforms. The InP platform benefits from monolithically integrated lasers and SOAs. It also integrates power monitors, passive components, and mode convertors for fiber attachments. The InP PICs are usually manufactured using 3″, 4″, or 6″ wafers and can involve multiple metal-organic chemical vapor deposition III−V material growths. On the other hand, silicon-based photonics or SiPICs use Si on insulator (SOI) to guide and manage light. It has been developed to be compatible with CMOS processing and can therefore take advantage of its mature technology and infrastructure to offer low-cost, high-volume manufacturable optics. The properties of the waveguiding Si layer in SOI can be tailored according to design and process requirements. Submicron Si epi thickness supports compact single-mode devices with tight bends to enable high-density PICs. These devices, however, can have higher waveguide losses and require mode expansion when coupling to fiber. SiPIC technologies with epi thickness above 1um suffer lower propagation loss but due to lack of sharp bends, they can face difficulties if larger scale integration is needed. Alternatively, light can be guided in silicon nitride (SiN) or other high-index insulators deposited on top of Si and silicon dioxide (SiO_2) to provide passive and low loss devices. Recently there have also been reports of photonics using bulk Si platform, where eposited polysilicon was used to guide light above the bulk Si [8]. The key building blocks of integrated SiPIC to enable optical transceiver for data communication are summarized below. For in depth discussion of each, refer to Chapters 2−5 for lasers, modulators, PDs, and passives, respectively.

7.2.1 Laser integration (external, flip-chip, heterogeneous)

Optical communication between Ethernet switches inside modern DCs uses SMFs, and the operating wavelengths of choice are within O-band, which spans 1260−1360 nm. III−V semiconductor lasers are needed to support these requirements of single-mode and O-band operation, and their design and performance are discussed in detail in Chapter 2. Once the lasers are fabricated, their light need to be coupled to the PIC waveguides. Since laser stability, noise and linewidth can degrade by optical back reflection, it is important to either use optical isolator between the laser and the PIC or lower the reflection of the PIC components to not impair the laser performance. In the latter case an isolator will still be needed at the Tx PIC to fiber interface to limit optical reflection from the fiber plant. Below are four common approaches used to integrate the laser in PIC-based optical transceivers.

1. Coupling through fiber, where a single or array of lasers are placed into an external package and the light is coupled through fiber into the PIC die. Depending on the design, the laser is actively temperature controlled using thermo-electric coolers (TECs) or passively cooled using air. The optical isolator is usually placed with the fiber attachment unit, which can include collimating lens optics. To protect the lasers from environmental conditions, such as humidity and electrostatic discharge over operating life, the lasers are often hermetically sealed in a pigtailed gold box. The pigtail must be polarization maintaining fiber (PMF) to preserve the TE polarization of the laser light. Fig. 7.5A shows the schematic of such a laser coupling scheme.
2. Coupling through lens and free space optics, where laser is placed on a submount next to the SiPIC die: Depending on the PIC design, the light can be coupled in through edge or surface couplers. In either case, the isolator is placed within the collimating lens system. One advantage of this approach is the elimination of the PMF pigtail and the associated cost. Another benefit is that this approach eliminates one coupling interface compared to the previous approach, so it has inherently lower coupling loss from laser to SiPIC. Fig. 7.5B shows the schematic of this laser assembly scheme.
3. Butt-coupling, where the laser is either attached directly or through a submount to bring the laser output facet to the PIC input facet. In this configuration, the laser and SiPIC waveguide modes must be designed to match to reduce the coupling loss. Often modes on both sides are expanded to improve alignment tolerance. Vertical alignment can be achieved using vertical references on both chips. The laser attachment is done through flip chip bonding using conventional bumping and die attach technologies. One big advantage of this approach is that it can be done at wafer level to enable full SiPIC performance SORT before any further assembly process. In this configuration the isolator is placed after SiPIC, and hence the SiPIC must be designed to keep the on-die reflection lower than laser feedback tolerance.
4. Coupling the light either heterogeneously or monolithically (for III−V PICs only) within the PIC platform. Unlike previous configurations, this approach does not assemble fully fabricated laser to SiPIC. Here, the III−V laser active region is either grown on a separate InP substrate or grown directly on silicon. If grown separately, pieces of the active material need to be heterogeneously bonded to the SiPIC. More recently there have also been successful demonstrations of directly growing quantum dot material directly on SOI wafers [9]. Once the active III−V material is integrated with silicon, wafer level CMOS

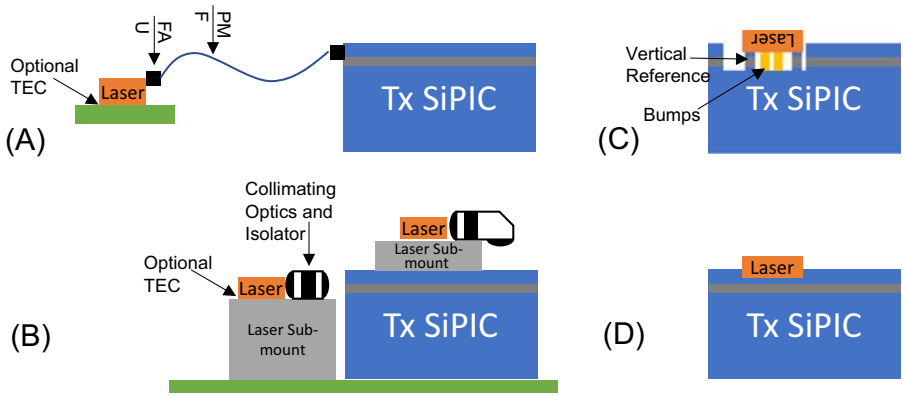

Figure 7.5 Various laser integration schemes: (A) coupling through fiber, (B) coupling using a submount, (C) coupling using butt coupling, and (D) coupling through heterogeneous integration.

fabrication processes are used to create the laser structure. In this implementation, the laser mode can be designed to overlap both Si and III−V to form a hybrid mode. Hence, light coupling to the Si waveguide can be enabled through integrated tapers that offer very low coupling loss. In addition, no hermetic sealing is required as the lasers are sealed within the same cladding that seals the rest of the SiPIC devices.

7.2.2 Modulators

Depending on application requirements, various modulator technologies have been demonstrated and integrated into PICs to realize transceiver modules for DC applications. Details on various modulator designs are covered in Chapter 3.

5. Electro-absorption modulators (EAMs) are commonly used in InP EML PICs. They are relatively compact and offer low capacitance and high speed. In addition, they do not require active control. However, their serviceable optical bandwidth is limited to tens of nanometers. Because they work based on light absorption, their maximum power capacity can be limited due to overheating effects. These devices have been implemented within SiPIC platform by depositing Ge/SiGe on Si. Franz−Keldysh (FK) type modulators work mostly over the C-band (1530−1565 nm) and L-band (1565−1625 nm) and provide high-bandwidth performance. The FK modulators, however, fail to deliver good performance at O-band. Instead, quantum-confined stark effect (QCSE) EAMs have delivered good high-bandwidth and low-voltage performance and have been integrated into SiPIC platforms.
6. Mach−Zehnder modulators (MZMs): MZMs are extensively deployed on SiPIC platforms. They operate based on free carrier plasma dispersion induced phase modulation in one or both arms of Mach−Zehnder interferometers (MZIs). Because the plasma dispersion effect in Si is relatively weak, the length of the phase section in each MZI arm is typically millimeters long. Delivering the high-speed signal to these long phase sections can be challenging and is done using either traveling wave electrodes or distributed driver scheme to ensure copropagation and velocity matching of the electrical and optical waves. High-speed Si MZMs are mature and have wide optical bandwidth, so they are the device

of choice for SiPIC-based products. Recently thin-film lithium niobate (TFLN) integration with SiPIC has demonstrated <2 V cm phase efficiency with <1-dB insertion loss [10,11]. Although the size of these devices can be large, they offer a path in developing low power modulators. It is noteworthy to also mention that integration of other materials, such as barium titanate ($BaTiO_3$) and lead zirconate titanate, has also been explored for achieving high-efficiency and low-loss modulators [12].

7. Microring modulators (MRMs): MRMs have recently emerged as a modulator of choice for large scale integrated photonic chips. These devices work based on the resonance effect, and optical modulation occurs by adjusting the phase of the resonance. Due to the relatively high-quality factor (Q) of MRMs, they can have high-modulation efficiency to lower the required drive voltage. However, high Q can compromise device bandwidth due to photon lifetime limitations and can cause nonlinearity in the device performance. Also, MRMs require tight control of the resonance and operating point, which are sensitive to bias, temperature, and environmental changes. These challenges have been addressed, and a full link demonstration using MRM-based PIC-passed DR4 Ethernet traffic [3]. The MRMs are additionally used in DWDM type applications, where multiple MRMs are designed within a single bus, each locked to and modulate a specific wavelength [13,14].

7.2.3 Receiver

The performance of the receiver circuitry for an optical interconnect is of paramount importance, as it determines the detection sensitivity of the communication link. Depending on the PIC design, the receiver SiPIC components can be integrated with the Tx components or can be designed as a separate photonic chip. To enable low-noise and high-signal contrast, the PD must have low dark current, low operating voltage and high responsivity across the operating temperature range. As the per-wavelength data rate climbs, the noise contribution of the PD becomes even more important. For photonic integration, planar waveguide PDs are needed and are typically based on germanium (Ge) grown on silicon with metal electrodes directly contacting the Ge or by forming lateral junctions in the silicon, to collect the photogenerated carriers. Using these same designs, charge multiplication layers can also be integrated to enhance the built-in field and increase carrier velocity to ballistic transport, thus causing an avalanche effect. These are called avalanche PDs and enable a balanced trade-off between high-gain and high-operating bandwidth that enable linear operation. These devices require applied reverse biases in the range of 5 V or more, so careful EIC and power delivery designs are needed to take advantage of these devices' high-gain bandwidth product. Please see Chapter 4 for a thorough review of high-speed integrated PDs.

7.2.4 Passives

Passive optical components are the basis of active optical components and the backbone of any successful PIC design. They need to be carefully optimized and fabricated to realize low overall link insertion loss. Passive components can be divided into two categories: input/output (*I*/*O*) devices and on-chip routing devices. In general, there are two major SiPIC passive platforms: single-mode submicron

waveguides and multimicrons MM waveguides. For *I/O*, both approaches can use either an edge coupler or a top-down coupler. Edge coupling is most efficiently done by mode conversion from a low index contrast material that enables a fiber-matched large mode diameter that is tapered down to the waveguide mode in the SOI layer. This edge coupler is typically realized in SiN. When integrated with V-grooves in the SiPIC, this coupler can enables passive fiber array assembly, which is a capability that has been demonstrated by many in the industry. Alternatively, near normal surface coupling can be realized using Bragg grating couplers (GCs) that allow for an almost 90 degrees turn of the waveguide light. While these devices do not require a dielectric for mode conversion, careful lithography and etch steps must be used to ensure low insertion loss. As photonic circuits have grown to add more and more functionalities, certain passive components, such as waveguide crossings and wide single-mode waveguides, have also been developed and refined to allow for low loss, flexible, and longer distance light propagation and routing. In addition, splitters, optical taps and wavelength multiplexers have been integrated and allow for even more advanced optical functions needed for next-generation optical interconnects. More information and details on passive components are covered in [13].

7.3 Pluggable transceivers for intradata center interconnects

Several photonics technologies in InP and Si including DMLs and EMLs, and SiPICs have been developed and commercialized for EPON/GPON fiber to the home and 40-, 100-, 200-, 400-, and 800-Gbps DC applications. In the following section, we will briefly describe current state of the art components and design trade-offs associated with each of these technologies as well as potential for scaling to higher bandwidth density solutions.

The optical transceiver modules consist of electronics, optical components, and light source, as shown in Fig. 7.6. They have been traditionally designed using discrete components and subassemblies. The transmitter optical subassembly (TOSA) uses either DMLs or EMLs in a separate laser module. The lasers are fiber coupled and are often assembled with the required drive electronics within the same module. On the receiver side, the receiver optical subassemblies also integrate individual or arrays of PDs with TIAs. The promise of silicon photonics and photonic integration is to enable lower cost scaling to higher data rates and channel counts, but it lacks a monolithically integrable light source. In addition, while SiPIC wafers are manufactured inside CMOS fabs, they do not necessarily integrate active CMOS circuits on the same wafer. These technology blocks therefore need to be assembled or integrated to form the full optical transceivers. Four different integration approaches are theoretically possible, and some have been demonstrated and deployed. The first approach integrates photonic components within the same die, while electronics and lasers are fabricated and assembled separately (Fig. 7.7A). The second

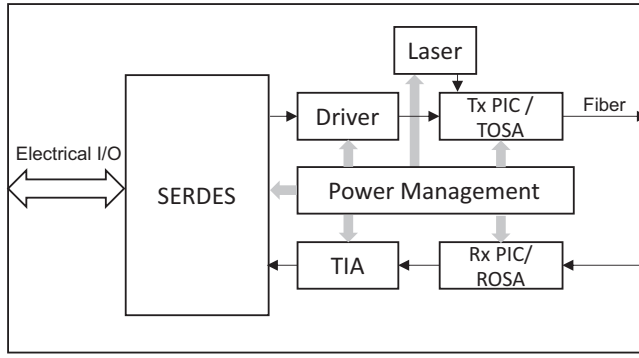

Figure 7.6 Block diagram of typical optical transceiver.

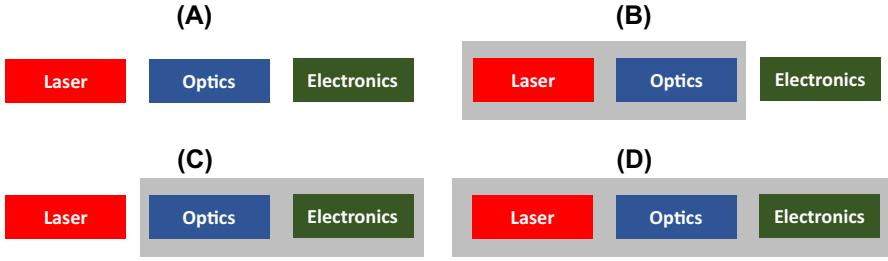

Figure 7.7 Various silicon photonic-integrated circuit integration configurations: (A) with discrete laser, optics and electronics, (B) with integrated laser and optics, (C) with integrated optics and electronics, and (D) with fully integrate laser, optics and electronics.

approach integrates the lasers within the photonic die, and the electronics are packaged separately (Fig. 7.7B). The third approach integrates the electronics with the photonics, and the laser is coupled externally (Fig. 7.7C). The last approach provides the ultimate integrated solution where electronics, PIC, and light source are all integrated together into the same chip (Fig. 7.7D). In this section, we first briefly introduce the available standards for transceivers and then discuss the transceiver technologies from discrete to the integrated photonics.

As DC size grows and becomes more complex, the required reach of the optical links also grows and BW increases. Different parts of this DC networking infrastructure have different requirements giving rise to the different transceiver types and specifications. Table 7.3 shows a comparison of the 100 GBase transceiver specifications. For endpoint connections with low BW (such as from TOR to NICs), parallel single-mode (PSM) type transceivers with only a single wavelength in each fiber are used. These PSM transceivers use an MPO connector where a fiber ribbon (with eight different fibers) is used for the interconnect. Splitter cables can be used to split the connections to four pairs of duplex connections for different NICs. Parallel optics-based transceivers also come in different rates 40 G, 100 G (PSM4)

Table 7.3 Comparing different types of 100 GBase transmitter specifications.

Description	**Value**			**Unit**
100 GBase	**DR4**	**FR4**	**LR4**	–
Electrical signaling rate, each lane (range)	26.5625 ± 100 ppm			GBd
Optical signaling rate per wavelength	53.125 ± 100 ppm			GBd
Modulation format	PAM4			–
Lane wavelength (range)	1304.5−1317.5	1264.5−1277.5 1284.5−1297.5 1304.5−1317.5 1324.5−1337.5	1294.53−1296.59 1299.02−1301.09 1303.54−1305.63 1308.09−1310.19	nm
Side-mode suppression ratio (SMSR), (min)	30			dB
Average launch power, each lane (max)	3	4.7	5.3	dBm
Average launch power, each lanea (min)	−5.1	−4.2	−3.4	dBm
Outer optical modulation amplitude (OMAouter), each lane (max)	2.8	4.5	5.1	dBm
Outer optical modulation amplitude (OMAouter), each lane (min)b	−3	−1.2	−0.4	dBm
Launch power in OMAouter minus TDECQ, each lane (min):	−4.4	−2.6	−1.8	dBm
For extinction ratio ≥ 4.5 dB	−4.4	−2.5	−1.7	dBm
For extinction ratio <4.5 dB				
Transmitter and dispersion eye closure for PAM4 (TDECQ), each lane (max)	3.4	3.3	3.4	dB
Average launch power of OFF transmitter, each lane (max)	−16	−30		dBm
Extinction ratio, each lane (min)	3.5			dB
RIN21.4OMA (max)	−132			dB/Hz
Optical return loss tolerance (max)	21.4	17.1	15.6	dB
Transmitter reflectance (max)	−26			dB

and 400 G (DR4) for different bandwidth requirements. For higher data traffic links, more data can be packed into a single fiber by using the coarse wavelength division multiplexed (CWDM) versions of the transceiver (CWDM4 for 100 G and FR4 for 400 G). Here, four wavelengths ranging from 1271 to 1331 nm are combined into a single waveguide, where each wavelength carries up to 100-Gbps data, giving a total per fiber bandwidth of up to 400 G. 100 GBase DR4 and FR4 transceivers are expected to have reach up to 2 km, while the LR4 transceivers are defined to carry 400 Gbps on a fiber pair up to 10 km.

To save power and cost, the lasers used in these transceivers are not thermally stabilized. Hence the lasing wavelengths will drift with temperature, and for WDM designs, a wavelength spacing of 20 nm is specified to accommodate the overall temperature drift without the need for temperature control. Depending on the transceiver type, Tx launch power, OMA, and optical return loss tolerance can vary, but side-mode suppression ratio (SMSR >30 dB), relative intensity noise (RIN < -132 dB/Hz), transmitter and dispersion eye closure quaternary penalty (TDECQ <3.4 dB), and ER (ER >3.5 dB) are relatively consistent [4]. For CWDM transceivers, a duplex LC optical connector is used, and only a single duplex fiber is needed for data transmission.

Next, we will review the three main technologies used for TOSA and Tx PIC, including DMLs, EMLs, and SiPICs.

7.3.1 Directly modulated lasers

DML is one of the oldest and simplest approaches to produce modulated optical signal by applying modulated current directly to the laser gain section. Traditional DMLs are fabricated on InP wafers and use buried heterostructure (Fig. 7.8A) or ridge technology to enable current and optical confinement. The power conversion efficiency of DMLs is quite high, and device compactness results in low manufacturing cost. Commercial optical transmitters based on DMLs are extensively deployed for short distance links with NRZ and PAM-4 modulation at 10 and 25 Gbaud. As the datacenter optical links advance to 50 Gbaud, the transient

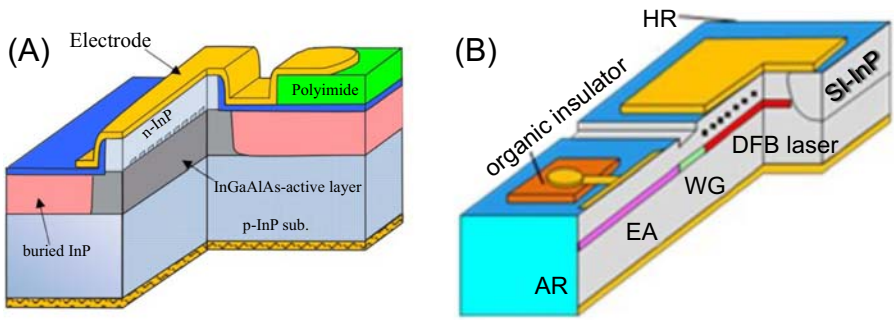

Figure 7.8 Cross sectional view of (A) directly modulated laser SR-LBH lasers [16] and (B) externally modulated laser electro-absorption/distributed feedback laser lasers [17].

chirp of DMLs limits the transmission distance for CWDM wavelength grid. Nevertheless, 2-km link at 106 Gbps PAM4 modulation was demonstrated with uncooled DMLs [15].

Furthermore, advanced research papers demonstrated drastic expansion of InP DML bandwidth to ~70 GHz by exploiting photon−photon resonance effect in compound cavities designed with passive feedback waveguide or distributed Bragg reflector mirror [18]. The photon-photon resonance effect produces resonant amplification of modulation sidebands by the presence of the side mode and is maximized when the laser is biased close to the mode hop condition, necessitating tight adaptive bias control over operating temperature and life of the device [19].

7.3.2 Externally modulated lasers

Integration of a single frequency continuous wave (CW) distributed feedback laser (DFB) with an external modulator enables superior scaling of modulation bandwidth, reduces chirp, provides ability to control output power and modulation characteristics independently, and typically enables improved performance over the wide operating temperature range required for pluggable modules.

High-efficiency electro-absorption (EA) and electro-optic MZMs in InP material system rely on the quantum confined Stark effect, which produces red shift of the multiquantum well absorption edge and changes the refractive index when reverse biased voltage is applied to p−i−n heterostructure. In InP technology, the integration of a DFB laser and a modulator is typically accomplished by selectively removing laser MQW stack and re-growing blue shifted MQWs stack for the modulator section. The performance of InP EA and MZMs is wavelength and temperature dependent and requires adjustment of the bias voltage to achieve the targeted ER and chirp. EA-DFBs are very compact and are widely deployed in 100 and 400 Gbps DR and FR transceivers in DCs, while InP Mach−Zehnder are typically integrated with widely tunable lasers designed for C-band operation and deployed in Metro and long-haul applications using NRZ and coherent modulation formats [20,21]. Several publications demonstrated scalability of the EA-DFB based optical transmitters, as shown in Fig. 7.8B, to 100 Gbps and even 200 Gbps operation [22]. It is expected that EA-DFB will be commercialized in 200 Gbps per lane front plate pluggable transceivers.

The need for four or eight DMLs or EA-DFBs to support CWDM wavelength grid architectures with 20-nm channel spacing prohibits fabrication of the arrays on one InP wafer due to detuning between operating wavelength and bandgap of the laser and modulator MQW stacks. Therefore four- or eight-channel optical modules implemented based on DMLs or EMLs rely on copackaging of discrete chips with additional optical components using several steps of high-precision alignment, and they frequently employ TEC to maintain the performance especially for higher baud rate applications at 50 Gbaud and beyond. Optical modules implemented using discrete chips have relatively large footprints that are still compatible with four- or eight-channel front plate pluggable transceiver form factors but are not scalable to bandwidth density required for CPO engines.

A very promising approach of expanding both performance and "integrability" of the DMLs and EMLs with passive optical components has been demonstrated using membrane technology in SiPIC [23]. In this technology, membranes of III−V MWQ active material are wafer bonded to SiPIC circuits and are used to fabricate CW lasers, DMLs or EMLs. Combining the benefits of high-optical confinement of the underlying SiO_2 layer and photon-photon resonance compound cavity designs, the researchers were able to achieve ∼60 GHz bandwidth [24].

7.3.3 Si photonics platform

Over the last decade silicon photonics has been rapidly penetrating datacenter and long-haul data transmission markets [25,26]. The success of Si photonics (SiPIC) technology is majorly determined by two factors: (1) high-refractive index contrast between Si and underlying SiO_2 layer leading to smaller waveguides and denser optical chips and (2) mature CMOS processing on large wafers, developed and funded by microelectronics industry, which drove down the cost of SiPICs while allowing for higher levels of integration and functionality. This disruptive technology has entered the market since 2011 starting with 4x14 Gbps transceivers [27]. After 2015, 100-G PSM4 type transceivers were introduced to DCs by multiple companies, including Intel, Luxtera and Mellanox [28−30]. Since then, many more SiPIC products, such as 100-G CWDM4, 100-G LR4, 200-G FR4, 400-G DR4 and 400-G FR4 transceivers, have been introduced, and millions of these transceivers are in the field [31]. So far, the SiPIC market has grown close to $1B and is expected to quadruple in five years [32]. This section will review the various SiPIC technologies and discuss details of a few notable ones.

Table 7.4 lists the mix of technologies developed and deployed in manufacturing of SiPIC transceivers including approaches for laser integration, key features of PIC, such as dimension of Si waveguides, types of modulators, optical coupling, and finally, solutions for copackaging and/or integration with EICs. In several cases, fully processed laser chips are flip-chip bonded on the SiPIC wafer with the laser mirrors formed either in the laser chips or in the SiPIC waveguides. This approach requires tight alignment between the optical mode of the laser chip and SiPIC waveguide and frequently results in parasitic reflections that can disturb the laser light-current, SMSR, and RIN performance.

An alternative technology for integrating an external laser at the wafer scale has been outlined in [53]. In this technology that was developed by Cisco (former Luxtera), an externally packaged light source called the LaMP, containing InP-based laser diode packaged with an optical isolator and a lens inside of a MEMS-based hermetic package, is built using wafer level processing and assembly (Fig. 7.9A). LaMPs are burned in and tested at wafer level prior to singulation to down-select known-good parts followed by wafer scale attachment of LaMPs to the SOI wafers containing photonics and electrical ICs (Fig. 7.9B). This wafer scale assembly technology is robust but has challenges for scaling to multiwavelength architectures.

Table 7.4 Various technologies developed using SiPIC for DC transceivers.

WL	Laser integration	PIC tech	Modulator	Fiber attach	EIC	References
O-Band	Heterogeneous	Submicron	TWMZM	Edge coupling	External and flip chip	[31,33–37]
O-Band	Heterogeneous	Submicron	MRM	Edge coupling	External	[3]
O-Band	Heterogeneous	Submicron	QCSE EAM		External	[38,39]
O-Band	Flip chipped	MM	QCSE EAM	Edge coupling	External	[40]
O-Band	Flip chipped	Submicron	TWMZM		External	[41]
C-Band	Flip chipped	MM	FK EAM	Edge coupling	External	[42]
O-Band	External	Submicron	TWMZM	Edge coupling	External	[43–46]
O-Band	Flip chipped	Submicron	MRM	Edge coupling	External	[47]
O-Band	Flip chipped	Submicron	TWMZM	GC	Integrated	[48,49]
O-Band	External	Submicron	TWMZM	GC	Integrated	[50–52]
O-Band	External	Submicron	TWMZM	Edge coupling	Flip chipped	[53,54]
O-Band	External	Submicron	TWMZM		External	[55,56]
O-Band	External	Submicron	MZM		External	[57]

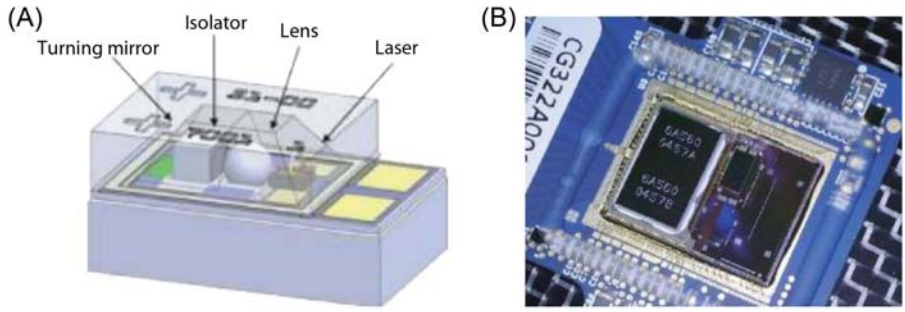

Figure 7.9 (A) Laser integration approach using LaMP and (B) fully assembled silicon photonic-integrated circuit with flip-chipped LaMP and electrical integrated circuit on top of silicon photonic-integrated circuit [53].

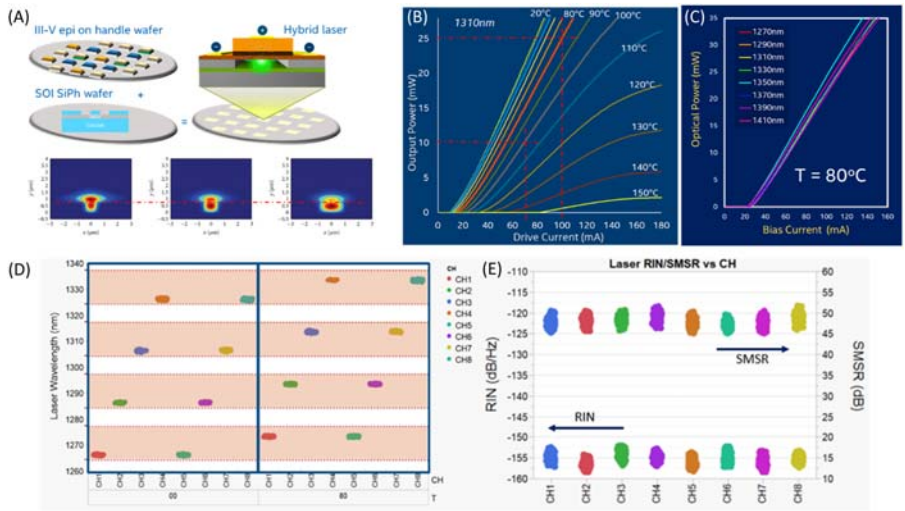

Figure 7.10 (A) Manufacturing flow of heterogeneously integrated laser and optical mode profile evolution as the light couples between the laser (left) and underlying Si waveguide (right). (B) Laser light coupled into Si waveguide as a function of laser bias current for a 1310-nm III/V-Si distributed feedback laser from 20°C to 150°C, in 10°C steps. (C) The light versus current characteristics of an eight-channel coarse wavelength division multiplexed transceiver fabricated with III/V-Si distributed feedback lasers showing uniformity of power across all channels at 80°C. (D) Laser wavelengths for a 2 × FR4 silicon photonic-integrated circuit at 0°C and 80°C. (E) Relative intensity noise and side-mode suppression ratio for all 2 × FR4 silicon photonic-integrated circuit channels.

Fig. 7.10 illustrates Intel's 300-mm heterogeneous integration process that combines high-performance multichannel laser arrays with high-baud rate SiPIC modulators, photodiodes, and passive optics on a flexible, scalable, low-cost manufacturing platform that supports a broad optical components portfolio [58].

The technology has been used to commercialize and ramp manufacturing of 100- and 400-Gbps optical transceivers, develop next-generation 800-Gbps solutions, and demonstrate high-density SiPICs for CPO engines and integrated optical *I/O*. The higher level of integration enabled by this heterogeneous technology will continue to drive penetration of SiPIC into next-generation data transmission as well as various other applications including sensing or neural networks for machine learning artificial intelligence. Fig. 7.10 also presents performance characteristics of 800-Gbps 2×400G FR4 SiPIC transmitter with eight heterogeneously integrated DFB lasers (1271, 1291, 1311, and 1331 nm), eight 100-Gbps traveling wave MZMs (TWMZMs), and two wavelength multiplexers (MUX). The data show good process control, leading to low variability in transmitter performance in terms of laser wavelength, SMSR, and RIN. The specifications for these three metrics are defined in Table 7.3. The wavelength ranges are indicated with the orange band in Fig. 7.10D, showing compliance across 80°C temperature range. In addition, for the data shown in Fig. 7.10E, the SMSR of greater than 30 dB and RIN of lower than −132 dB/Hz are required by the specification.

Other approaches to heterogeneous integrate III−V material on SiPIC are described in [36,38]. In the technology presented in [38], the unprocessed III−V chiplets are metal-bonded directly to the Si substrate in the trenches etched into SOI wafers. The side gaps between the III-V material and the SOI material are filled with SiO_2 and α-Si to ensure continuous optical confinement. The vertical waveguide alignment is controlled by the material layer thickness and selective etching, while the lateral waveguide alignment is ensured by photolithography. The integrated III−V multiquantum-well material is used to fabricate DBR lasers and EA modulators enabling development of integrated optical transmitters.

Majority of the qualified SiPIC based products in the market are based on MZM, these devices are often integrated with drivers through wire bonding. More recently and as the baud rate increases, there have been products that use flip chip bonded EIC on the SiPIC that are connected to the host with through-silicon vias (Fig. 7.9B). Such assembly technique enables testing the full SiPIC subsystem at wafer level that enables lower manufacturing cost [53]. Alternatively there have been other solutions that integrate the BiCMOS technology with the photonic process and deliver integrated EIC with SiPIC [48,49]. In addition, MRMs and EAMs have been gaining momentum to address the increasing bandwidth density demand [3,40,47].

Silicon photonics technology has been deployed in several generations of Ethernet transceivers starting at 40 Gbps and advancing to 800 Gbps. Intel 300 mm heterogeneous SiPIC platform with integrated multichannel laser arrays, high-speed modulators, photodiodes, and multiplicity of passives components enabled rapid development. Product ramps to high volume for several generations of pluggable Ethernet transceivers have delivered eightfold increase of throughput from 100 to 800 Gbps in the last 6 years since 2016 [33−35]. Fig. 7.11 schematically illustrates the architectures employed by Intel to scale along the three technology vectors of wavelengths/fiber, baud rate, and modulation format.

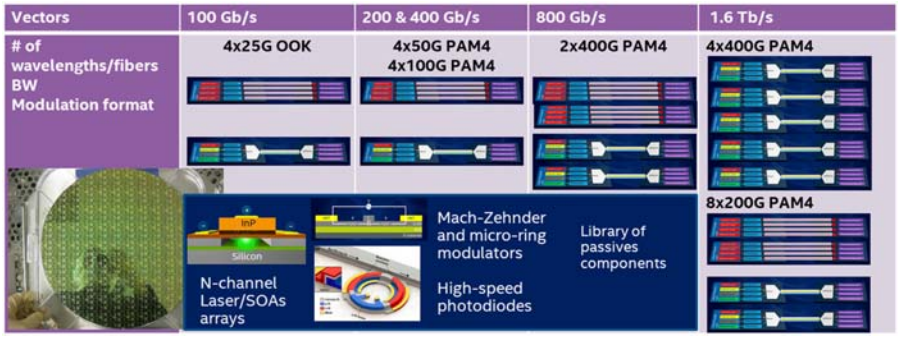

Figure 7.11 Architectures options to scale along the three technology vectors of wavelengths/fiber, baud rate, and modulation format.

7.4 Photonic integration/copackaging for data center interconnects

Front-plate pluggable optical modules have scaled in channel count and line rate to meet the data throughput demand of 25.6 T (32×800 G) switches and are poised to deliver solutions for 51.2 T (32×1.6 T) switches. It is generally accepted that scaling to future high-bandwidth switching nodes is becoming more challenging due to high loss in the electrical interconnects between the switch ASIC and front plate pluggable modules. As discussed in Chapter 1, the solution to reduce the electrical power consumption at higher data rate is to shorten the electrical channel by bringing the optics closer to the switch ASIC. For CPO interconnects to replace the front plate pluggable optical modules, they must have dense integration of all needed transceiver components so they could be placed onto the same package as the switch ASIC [4,59,60,61]. This tight packaging scheme enables significant SERDES power reduction as it shrinks the electrical channel between switch and optics from many inches down to tens of millimeters in length and alleviates front panel physical congestion. This power saving can be critical in datacenters in order to accommodate the next generation of switch technology within existing power envelopes and infrastructures [62]. However, the CPO optical module architecture imposes stringent requirements on optical components including bandwidth density, power conversion efficiency, serviceability and reliability.

High-bandwidth switches have an electrical bandwidth escape density of >200 Gbps/mm for a linear beachfront. To avoid package size increase, the PIC design must enable similar optical bandwidth density or accommodate reasonable fan-out of the electrical channel within the package. Cutting-edge PICs should also be able to accommodate high-bandwidth SI impairments, dynamically tune thermally sensitive devices, employ optical monitors to ensure health of the optical signal and optimal operation of the components, and handle polarization diversity. In addition, if the design integrates on-chip III−V, then the PIC must provide sufficiently high-optical power or gain for a healthy link budget margin.

As the number of on-chip components in a CPO engine increase, more stringent requirements are placed on tighter footprint of these components. MZMs, which typically consume 2−3 mm of linear waveguide length, arrayed waveguide grating multiplexers, polarization splitter-rotators are examples of popular devices that consume a considerable amount of area. Next-generation CPO circuits employ MRMs, cascaded MZI-based MUXs and DeMUXs, and connector-based polarization handling components to further reduce the chip size. High-efficiency optical coupling, low-loss waveguides, and advanced back-end metal schemes enable the integration of the aforementioned functionalities into the PIC. These innovations enable a power-efficient link that helps alleviate the overall power consumption of next-generation DC switch *I/O* communication.

Integrated circuits offer the advantage that all components are fabricated at once and are connected to each other without excess interface loss. By leveraging the economies of scale for manufacturing and testing at the wafer level, SiPICs used for CPO stand to undercut costs of assembly that are incurred for traditional pluggable modules. In addition, the ideal photonic components technology needs to be able to sustain rapid volume ramps to millions of units per year as well as scaling of the architectures and performance for several generations of transceivers to come.

The optical transceiver or PIC has a host-side electrical interface and a line-side optical fiber interface. The electrical and optical packaging choices for CPO engines impact system integration, and the fiber connection and method therein also drive key downstream packaging choices. The use of edge couplers enabled by SiN mode converters and silicon v-grooves have been a staple of the silicon photonics market [3] for many years. This approach enables a broadband optical solution, amenable for CWDM wavelength requirements and allows on-chip components, such as polarization rotator-splitter, to handle incoming polarization variation. The edge-coupled solution enables a linear beachfront optical *I/O*, which must support the required total bandwidth as discussed previously. Typically, v-grooves are built on a 250-μm pitch (matching SMF fiber ribbon pitch) with more dense pitch designs (such as 127 μm) being discussed today. The v-groove fiber attachment is challenging because it is a permanent attachment that requires epoxy that limits the maximum temperature allowed for subsequent assembly steps. Also, edge coupling complicates testing, particularly functional testing at the wafer level. So, GCs are alternatively integrated to enable simpler testing. Another challenge is the integration of SiN, which presents film stress related mechanical challenges for integration of key devices and features, such as through-vias that can greatly improve electrical *I/O* density into the chip. In addition, v-groove based assembly requires connectorized short fibers, referred to as a fiber pigtail, which increase packaging complexity because of the manual or semiautomated handling of such parts.

Top-down optical *I/O* offers the ability to scale bandwidth escape in two dimensions, improving the electrical to optical bandwidth density. This transition from edge coupling to top-down surface coupling is akin to the wire-bond to solder bump transition of the CMOS industry nearly a decade ago. Advances in fabrication processes make GCs a viable technology for near-normal optical *I/O* [63], and other

approaches using angled mirrors and lens integration [64] can offer more broadband performance. GCs do not rely on mode conversion of nitride materials. However, their narrower optical bandwidth and polarization dependence require careful design in the overall PIC and optical link. Top-down *I/O* allows for socket-based fiber attachment that can eliminate epoxy and enable functional testing of the final optical facets [65]. The resulting removable connector can also eliminate the fiber pigtail and enables a PIC that can be handled by existing, fully automated handling tools at high-volume assembly vendors.

The electrical interface of the CPO engine to the host chip can be greatly influenced by the optical fiber attach solutions outlined above. Nevertheless, tighter integration of the optics and host ASIC is preferred and is possible by placing both on a single package substrate, which is known as multichip module. The electrical *I/O* density here is increased by $3-5\times$ as compared to the grid array embodiment using microbumps. As stated previously, this tighter packaging requires the same testing capabilities and reliability requirements for CMOS components. One step further is also possible by placing the host ASIC and CPO engine on the same interposer, either using microbump for very dense electrical bandwidth routing or socketed engines to address the serviceability. If the engine cannot be screened prior to commitment to the final package, an electrical socket solution is often preferred. Moving one more step further, the photonic engine can alternatively be placed on a nearby socketed package and routing of electrical signals is done via PCB. This is referred to as near-package optics. In this case, the electrical *I/O* density is typically a grid array with a ball pitch of 0.3 mm or more.

These various form-factors present opportunities to reduce power consumption from the electrical serializer/deserialize (SerDes) of the host ASIC to the analog front-end of the photonic components. As the SI improves by going from routing on PCB to substrate to interposer, the electrical link between the host ASIC and analog front-end circuits can improve from a few pico-Joules per bit to below 1 pJ/bit. However, such a reduction in power implies the SerDes are optimized for use with a specific photonics solution, which has product-level implications for up-front commitment to a particular optical *I/O* solution. In recent years, there has been an increase in adoption of a direct-drive optical *I/O* solution. This is defined as an electrical link between the host ASIC and photonics front-end circuits that does not include a clock-data recovery and/or re-timer component. Direct drive can reduce power and improve overall latency of the overall *I/O* solution and must be able to equalize for the bandwidth limitations faced by the optical components. Inherently, direct-drive schemes have superior power efficiency for multichip and copackaged solutions, where the SI is much better.

An example of fully functional CPO engine compliant to 400-G DR4 specifications was presented in [3,66] and is shown in Fig. 7.12, demonstrating the power of integration. The copackaged SiPIC with its fiber I/O exceeds the bandwidth of 100-G pluggable modules by $16\times$ while improving the bandwidth density over 40 times and reducing the power consumption by 30%. This CPO engine was also integrated within a platform with a 12.8-Tbps Ethernet switch, as shown in Fig. 7.13 (top left image) [59]. More recently there have been more concept and operational

Figure 7.12 (A) Comparison of 100-Gbps quad small form-factor pluggable package versus 1.6-Tbps copackaged optics engines, (B) inside of a copackaged optics engine using fully integrated silicon photonic-integrated circuits with heterogeneously integrated lasers.

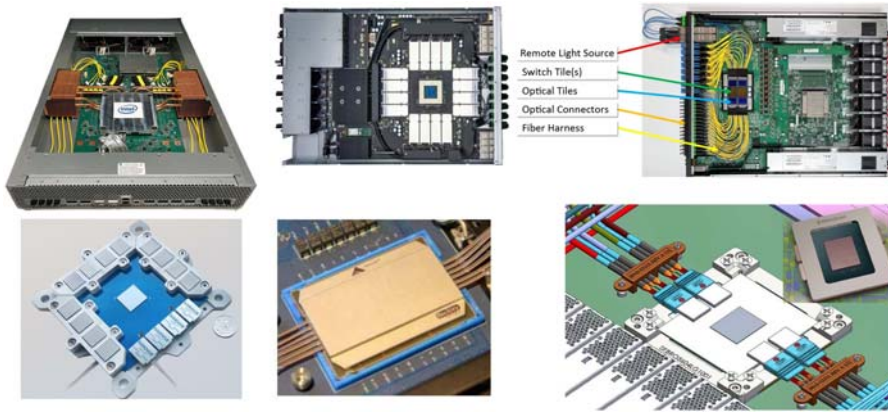

Figure 7.13 Various implementations of copackaged optics across industry.

demonstrations of CPO across the industry, as shown in Fig. 7.13 [59–61,67,68]. All of these demonstrations use SiPICs. Other than the demonstration by Intel, the other demos use an external laser source (ELS) that is coupled to the SiPIC using PMFs. These ELS sources are designed within a pluggable form factor to allow for serviceability. A mix of MRM and TWMZM have been used for the modulation scheme for these systems. Also, on the fiber attach side, both edge coupling using V-groove and surface coupling have been utilized in these demonstrations.

In the end, it is important to note that most of the work on CPO has targeted IMDD Intra-DC switching applications and focused on SiPIC engines, there have also been other CPO demonstrations for DCI coherent using SiPIC [17] as well as using legacy VCSEL technology for shorter reach intra-DC applications [69] and [70].

7.5 Future directions

7.5.1 Scaling datacenters to 102-Tbps nodes and beyond

The 51.2-Tbps Ethernet switch generation is poised to have 512 electrical and corresponding optical 100-Gbps PAM4 lanes. Doubling the BW required for the subsequent 102.4-Tbps switch node and beyond can be achieved by scaling lane speed, bits per symbol, number of lanes, or a combination of these.

When the data rate increases, the channel loss between the switch and optical module increases exponentially, which can result in higher energy consumption per bit. One option to avoid doubling the line rate is to use PAM6 (or higher order modulation) to increase the BW without increasing the fundamental frequency (i.e., keeping the same loss). Although such solutions may be achievable in the electrical domain where the noise is low, it is challenging in the optical domain due to higher noise levels. One potential solution is for the electrical lane to use PAM6 modulation while the optical modulation stays at PAM4. In this case, gearboxes are needed for the conversion, which again increases the overall power consumption.

Moving to 112 Gbaud electrical baud rate (224 Gbps PAM4) limits the electrical reach if targeting the same power envelope. One solution is to move the optical modules closer to the data switch die (i.e., using CPO), or to use more re-timers on the electrical channel. Reducing the channel between the optical module and the switch will eventually change the electrical interconnect. At 100 Gbps such a proposal can enable direct drive, which uses only one SERDES in the link, that is, at the switch end. It requires the optical module to be close enough to the switch that the SERDES inside the optical module can be removed. The short electrical interconnect still allows the design of a the SERDES to support the channel loss and meet the power consumption budget. On the other hand, at 224-Gbps PAM4 pluggable optics would still be possible by using lower loss materials (higher cost) in the boards and/or using multiple retimers (higher power). In either case at the 102.4-Tbps node, rearchitecting the switch platform is needed to enable higher electrical power delivery, higher power efficiency, and better cooling.

Once the optical module is very close to the switch die and the number of electrical lanes can scale within the package, the solution for lowest power consumption is no longer higher modulation scheme. Staying with NRZ and perhaps even lowering the data rate to make the link massively parallel could allow the overall power consumption on the electrical side to be lowered. This will shift the problem from one of power efficiency to also edge bandwidth density or shoreline. The switch dies have limited shoreline, so the IOs would need to be extended deeper into the die.

To further improve the efficiency and density of the shoreline, microbump interconnect could be developed. Smaller bump pitch and size increase the interconnect density and reduce interconnect capacitance, respectively. In addition, shorter distance enables very little power IOs, as only one low power SERDES, at the receiver end of the link, is needed to clean up and reconstruct the data stream.

In the future, hybrid bonding could be used to enable further high-density, 3D assembly integration of ICs. For example, die-stack the PIC onto the switch die directly.

7.5.2 Photonic-integrated circuit scaling to higher data rates

The scaling of the PIC is another approach to consider. The transceiver optical data rate per fiber can scale along different axes, as can be seen in Fig. 7.14. The modulation speed, modulation order, number of wavelengths and number of polarizations can all be used to scale the data rate of an optical system.

The current state of the art for optical modules uses a maximum of 4 wavelengths (O-band CWDM), a single polarization, and 112-Gbps PAM4 optical line rate per fiber (meeting IEEE 802.3-bs standards). This standard will also govern the 51.2-Tbps switch node. Such an optical data rate for the 51.2-Tbps node has been made possible through data rate doubling on the electrical side and fiber count doubling on the optical side. As discussed in the previous section, at 102.4-Tbps switch node the optics will likely stay with PAM4 signaling to avoid higher level modulation scheme's increase in SNR penalty. The actual optical line rate doubling to 224-Gbps PAM4 is the most likely path to enable the next node. There have been several demonstrations of modulators and detectors at this speed. Another option is to double the optical channel count by doubling the number of wavelengths. The issues here would be the wavelength spacing as well as the need to double the number of electrical drive signals. Scaling with coarse division multiplexing (20-nm spacing) would still ensure an athermal solution, but at 100-Gbps PAM4, the wide

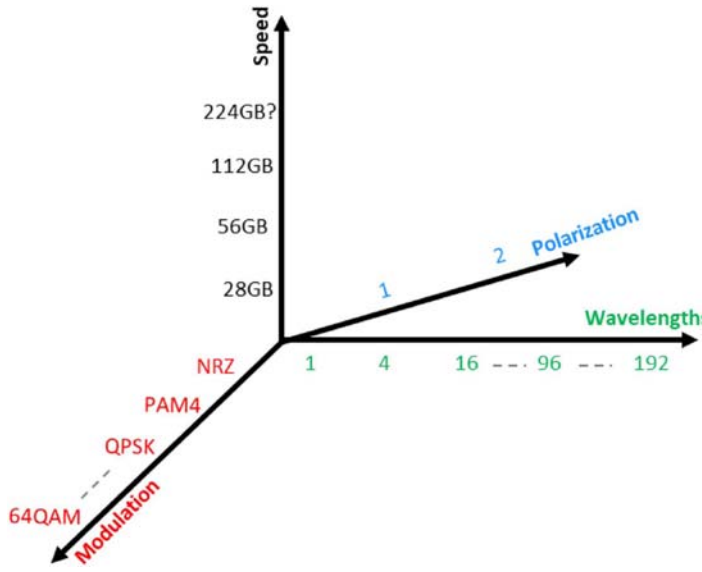

Figure 7.14 Data rate scaling vectors for future photonic-integrated circuit generations.

wavelength span will limit the reach of the optical link due to dispersion at outer channels far from 1310 nm. Alternatively, one can keep the same wavelength span but reduce the channel pitch to 10 nm. Doing so would require either temperature control of the PIC or limiting the operating temperature of the PIC. In addition, new MUX and DeMUX schemes would need to be developed.

The coherent transmission formats can enable further scaling to 400 and 800 Gbps per lane and beyond by employing higher order modulation and adding polarization and phase as additional scaling paths while delivering higher link budgets and maintaining longer reach (10 km and beyond) capability. Traditionally, increased complexity and cost of the coherent photonic components and higher power consuming DSPs make them less attractive relative to IMDD solutions for the cost sensitive SR applications, but the cost and power differences are being significantly reduced since the large number of IMDD lanes can be replaced with just one coherent channel. The coherent transmission is expected to become more relevant to short distance applications and may displace IMDD formats even for intra DC applications for 200-T switch generation. More details of the coherent transmission for DC connectivity are discussed in Chapter 6.

References

[1] 12 Largest Data Center in the World. <https://www.rankred.com/largest-data-centers-in-the-world/>, 2022.
[2] Cisco Global Cloud Index, 2016−21.
[3] S. Fathololoumi, et al., 1.6 Tbps silicon photonics integrated circuit and 800 Gbps photonic engine for switch co-packaging demonstration, J. Light. Technol. 39 (4) (2021) 1155−1161.
[4] IEEE Standard for Ethernet—Amendment 10: Media Access Control Parameters, Physical Layers, and Management Parameters for 200 Gb/s and 400 Gb/s Operation, IEEE Std 802.3bs-2017.
[5] Merits of selecting CWDM for 10km SMF PMD—IEEE 802. (2008). <https://www.ieee802.org/3/ba/public/mar08/traverso_01_0308.pdf>.
[6] Direct Drive broadcom. (2022) <https://www.broadcom.com/info/optics/cpo#:~:text=Co%2DPackaged%20Optics%20(CPO)%20is%20an%20advanced%20heterogeneous%20integration,generation%20bandwidth%20and%20power%20challenges>.
[7] OIF Hot Topic Fact Sheet—Common Electrical I/O (CEI)-112. (2019). <https://www.oiforum.com/wp-content/uploads/OIF-Hot-Topic-Fact-Sheet_CEI_FINAL.pdf>.
[8] A. Atabaki, et al., Monolithic optical transceivers in 65 nm bulk CMOS, in: Proceedings of the Optical Fiber Communications Conference and Exposition (OFC), 2018, pp. 1−3.
[9] C. Shang, et al., Low-threshold epitaxially grown 1.3-μm InAs quantum dot lasers on patterned (001) Si, IEEE J. Sel. Top. Quantum Electron. 25 (6) (2019) 1−7.
[10] H. Li, et al., High performance thin-film lithium niobate MZ modulator ready for massive production, in: Proceedings of the Optical Fiber Communications Conference and Exhibition (OFC), 2022.

[11] K. Luke, P. Kharel, C. Reimer, L. He, M. Loncar, M. Zhang, Wafer-scale low-loss lithium niobate photonic integrated circuits, Opt. Express 28 (2020) 24452−24458.
[12] A. Rahim, et al., Taking silicon photonics modulators to a higher performance level: state-of-the-art and a review of new technologies, Adv. Photon. 3 (2) (2021) 024003.
[13] S. Fathololoumi, et al., Highly integrated 4 Tbps silicon photonic IC for compute fabric connectivity, in: IEEE Hot Interconnect, 2022.
[14] X. Zheng, A.V. Krishnamoorthy, Si photonics technology for future optical interconnection, in: X. Liu, E. Ciaramella, N. Wada, and N. Chi (Eds.), Optical Transmission Systems, Subsystems, and Technologies, vol. 8309 of Proceedings of SPIE, Optica Publishing Group, 2011, paper 83091 V.
[15] T. Nakajima, et al., 106-Gb/s PAM4 operation of directly modulated DFB lasers from 25 to 70°C for transmission over 2-km SMF in the CWDM range, J. Light. Technol. 40 (6) (2022) 1815−1820.
[16] K. Nakahara et al., 112-Gb/s PAM-4 uncooled directly modulated BH lasers, ECOC, 2021.
[17] B. Buscaino, B.D. Taylor, J.M. Kahn, Multi-Tb/s-per-fiber coherent co-packaged optical interfaces for data center switches, J. Light. Technol. 37 (13) (2019) 3401−3412.
[18] Y. Matsui, et al., Isolator-free > 67-GHz bandwidth DFB + R laser with suppressed chirp, in: OFC 2020, Th4A.1.
[19] D. Che, et al., Long-term reliable >200-Gb/s directly modulated lasers with 800GbE-compliant DSP, in: OFC 2021.
[20] Y. Akulova, InP photonic integrated circuits for high efficiency pluggable optical interfaces, in: OFC 2015, W3H.1.
[21] Y. Nakai, et al., Uncooled operation of 53-GBd PAM4 (106-Gb/s) EA/DFB lasers with extremely low drive voltage with 0.9 Vpp, J. Light. Technol. 37 (7) (2019) 1658−1662.
[22] Y. Matsui, et al., 224-Gb/s PAM4 uncooled operation of lumped-electrode EA-DFB lasers with 2-km transmission for 800GbE application, in: OFC 2021, Tu1D.1.
[23] T. Aihara, et al., Lateral current injection membrane buried heterostructure lasers integrated on 200-nm-thick Si waveguide, in: OFC 2018, W3F.4.
[24] N.-P. Diamantopoulos, et al., 60 GHz bandwidth directly modulated membrane III-V Lasers on SiO2/Si, J. Light. Technol. (2022).
[25] H. Wang, et al., Silicon photonic transceivers for application in data centers, J. Semicond. 41 (2020) 101301.
[26] Z. Zhou, et al., Development trends in silicon photonics for data centers, Opt. Fiber Technol. V 44 (2018).
[27] <https://www.lightwaveonline.com/optical-tech/article/16660502/luxtera-introduces-silicon-cmos-photonicsbased-optical-engine> 2011.
[28] <https://www.businesswire.com/news/home/20150323005646/en/Luxtera-Debuts-1310nm-100G-PSM4-QSFP28-Module-and-Silicon-Photonics-Chipset-at-OFC-2015> 2015.
[29] <https://www.datacenterknowledge.com/archives/2016/08/17/intels-first-silicon-photonics-gear-for-data-center-networks-now-shipping> 2016.
[30] <https://www.businesswire.com/news/home/20170906005334/en/Mellanox-and-Accelink-Partner-to-Provide-100Gbs-PSM4-Ethernet-Transceivers> 2017.
[31] R. Blum, Integrated silicon photonics for high-volume data center applications, in: Proc. SPIE 11286, Optical Interconnects XX, 112860M, 28 February 2020.
[32] <https://www.mordorintelligence.com/industry-reports/silicon-photonics-market> 2022.

[33] H. Yu, et al., 100Gbps CWDM4 silicon photonics transmitter for 5G applications, in: OFC 2020, W3E.4.
[34] H. Yu, et al., 400Gbps fully integrated DR4 silicon photonics transmitter for data center applications, in: OFC 2020, T3H.6.
[35] H. Yu, et al., 800 Gbps fully integrated silicon photonics transmitter for data center applications, in: OFC 2022, M2D.7.
[36] S. Menenzo, (Integrated or not?) Laser source for a few pJ/bit DWDM links, in: Optical Fiber Conference, 2022.
[37] G. Li, et al., 100 Gb/s CWDM transmitter and receiver chips on a monolithic Si-photonics platform, in: Proceedings of the IEEE 13th International Conference on Group IV Photonics (GFP), 2016, pp. 164−165, doi:10.1109/GROUP4.2016
[38] G. Fish, A. Fang, Heterogeneous integration: extending the functionality of silicon photonics with III-V materials, in: Proceedings of the European Conference on Optical Communication (ECOC), 2015.
[39] T.N. Huynh, et al., Four-channel WDM transmitter with heterogeneously integrated III-V/Si photonics and low power 32 nm CMOS drivers, J. Light. Technol. 34 (13) (2016) 3131−3137.
[40] C. Minkenberg, et al., Reimagining datacenter topologies with integrated silicon photonics, J. Opt. Commun. Netw. 10 (7) (2018) 126−139.
[41] H. Zhu, S. Anderson, N. Karfelt, L. Jiang, Y. Li, R. Boeck, et al., Low-cost 400 Gbps DR4 silicon photonics transmitter for short-reach datacenter application, Nanomaterials (2021).
[42] M. Asghari, et al., Silicon photonics transmitters and receivers for 4x25 Gb/s interconnects, in: CLEO, 2014.
[43] Y. Ma, et al., An all-silicon transmitter with co-designed modulator and DC-coupled driver, in: Proceedings of the Optical Fiber Communication Conference (OFC), 2019.
[44] C. Xie et al., Real-time demonstration of silicon-photonics-based QSFP-DD 400GBASE-DR4 transceivers for datacenter applications, in: Proceedings of the Optical Fiber Communications Conference and Exhibition (OFC), 2020.
[45] L. Soldano, et al., Multi-wavelength 100Gb/s silicon photonics based transceiver with silica mux/demux and MEMS-coupled InP lasers, in: Proceedings of the Optical Fiber Communication Conference, 2017.
[46] Y. Zhao, et al., Silicon photonic based stacked die assembly toward 4×200-Gbit/s short-reach transmission, J. Light. Technol. 40 (5) (2022) 1369−1374.
[47] D.F. Logan, et al., 800 Gb/s silicon photonic transmitter for copackaged optics, in: Proceedings of the IEEE Photonics Conference (IPC), 2020.
[48] S. Meister, et al., Silicon photonics for 100 Gbit/s intra-data center optical interconnects, in: Proceedings of the SPIE 9753, Optical Interconnects XVI, 975308, 2016.
[49] S. Lischke, et al., Advanced photonic BiCMOS technology with high-performance Ge photo detectors, in: Proceedings of the SPIE 11088, Optical Sensing, Imaging, and Photon Counting: From X-Rays to THz 2019, 110880M.
[50] C. Gunn, CMOS photonics for high-speed interconnects, IEEE Micro 26 (2) (2006) 58−66.
[51] T. Pinguet et al., 25 Gb/s silicon photonic transceivers, in: Proceedings of the Nineth International Conference on Group IV Photonics (GFP), 2012.
[52] G. Masini, et al., Silicon photonics transceivers for high-speed data communication, in: 2017 Conference on Lasers and Electro-Optics Pacific Rim. A. Mekis, et al., High-speed silicon photonics optical transceivers, in: Proceedings of the IEEE Summer Topicals Meeting Series, MC3.1, 2017.

[53] T. Pinguet, et al., High-volume manufacturing platform for silicon photonics, Proc. IEEE 106 (12) (2018).
[54] F. Boeuf, et al., Silicon photonics R&D and manufacturing on 300-mm wafer platform, J. Light. Technol. 34 (2) (2016) 286−295.
[55] N.M. Fahrenkopf, et al., The AIM photonics MPW: a highly accessible cutting edge technology for rapid prototyping of photonic integrated circuits, IEEE J. Sel. Top. Quantum Electron. 25 (5) (2019) 1−6.
[56] E. Timurdogan et al., 400G silicon photonics integrated circuit transceiver chipsets for CPO, OBO, and pluggable modules, in: OFC 2020, T3H.2.
[57] K. Shastri, CMOS photonics, in: Proceedings of the Asia Communications and Photonics Conference and Exhibition (ACP), 2009.
[58] Park, et al., InP photonic integrated circuits for high efficiency pluggable optical interfaces, IEEE J. Sel. Top. Quantum Electron. 12 (6) (2006) 1657−1663.
[59] <https://newsroom.intel.com/news/intel-demonstrates-industry-first-co-packaged-optics-ethernet-switch/#gs.y8d38z> 2020.
[60] <https://ranovus.com/ranovus-announces-single-chip-odin-analog-drive-for-its-second-generation-co-packaged-optics-in-hyperscale-data-center/> 2021.
[61] B. Welch, Co-packaged optics integration, in: Proceedings of the EPIC Online Technology Meeting on CPO, 2020.
[62] R. Stone et al., Co-packaged optics for data center switching, in: Proceedings of the European Conference on Optical Communications (ECOC), 2020.
[63] Y. Luo, Z. Nong, S. Gao, H. Huang, Y. Zhu, L. Liu, et al., Low-loss two-dimensional silicon photonic grating coupler with a backside metal mirror, Opt. Lett. 43 (2018) 474−477.
[64] A. Israel, F. Ulfan, L. Pascar, S. Menezo, Q. Wilmart, S. Malhouitre, et al., Photonic plug for scalable silicon photonics packaging, in: Proceedings of the SPIE 11286, Optical Interconnects XX, 1128607, 2020.
[65] A. Janta-Polczynski, E. Cyr, J. Bougie, A. Drouin, R. Langlois, D. Childers et al., Towards co-packaging of photonics and microelectronics in existing manufacturing facilities, in: Proc. SPIE 10538, Optical Interconnects XVIII, 105380B, 2018.
[66] R. Mahajan, et al., Co-packaged photonics for high performance computing: status, challenges and opportunities, J. Light. Technol. 40 (2) (2022) 379−392.
[67] <https://ranovus.com/ranovus-launches-its-single-chip-odin-silicon-photonic-engine-to-support-mL-ai-workloads-for-data-center-and-5g-mobility/> 2020.
[68] <http://www.ragilenetworks.com/company/news/21> 2021.
[69] P. Maniotis, L. Schares, D.M. Kuchta, B. Karacali, Improving data center network locality w/ co-packaged optics, in: Proceedings of the European Conference on Optical Communication (ECOC), 2021.
[70] L. Dong, et al., Densely packed 1.1 μm band vertical cavity surface emitting laser array for co-packaged optics, Jpn. J. Appl. Phys. 61 (2022) SK1011.

This page intentionally left blank

Photonic switch fabrics in data center/high-performance computing networks

8

Qixiang Cheng[1], Chunhui Yao[1], Nicola Calabretta[2], Ripalta Stabile[2], Keijiro Suzuki[3], Hitoshi Kawashima[3], Weijie Tang[4], Madeleine Glick[5], Tao Chu[4], Kazuhiro Ikeda[3], Ryosuke Matsumoto[3], Shu Namiki[3], Keren Bergman[5] and Richard Penty[1]

[1]University of Cambridge, Cambridge, United Kingdom, [2]Eindhoven University of Technology, Eindhoven, The Netherlands, [3]The National Institute of Advanced Industrial Science and Technology, Tsukuba City, Japan, [4]Hangzhou, Zhejiang Province, Zhejiang University, P.R. China, [5]Columbia University, New York, NY, United States

8.1 Introduction

As introduced in Chapter 1, the world has witnessed the exponential growth of data traffic (doubling almost every 2 years) in warehouse scale data centers and large-scale high-performance computing (HPC) systems in the past few decades [1]. Due to the increasing demands for various data intensive applications, such as video services, AI and machine learning, and cloud computing, the data center network (DCN) had entered the so-called "Zettabyte (ZB) Era" with annual global traffic of 1.5 ZB in 2017 and this is expected to reach over 4.8 ZB per year in 2022 [2,3]. To serve the increasing traffic demands, the current trend has been driven towards mega-data centers with hundreds of thousands of servers for the benefits of economy of scale [4,5]. On the other hand, the traffic flow of the data center/HPC networks has shifted from North−South to East−West, that is, the internal data traffic between servers and storages surpasses inbound or outbound data traffic [6]. These trends pose unprecedented challenges to the intradata center interconnects in terms of bandwidth and connectivity.

In data center/HPC networks, in addition to links and network interfaces that provide point-to-point data communication, switches play an essential role as they guarantee the appropriate connections between input links to target links [7]. In the face of increasing bandwidth demands, conventional electrical transmission suffers from limited bandwidth (usually tens of Gb/s) on a single wire, making it challenging to achieve the data rate of hundreds of Gb/s as required by modern data centers and HPCs [8]. In addition to the link-level limitations, electrical switches also encounter bottlenecks due to not only the slowing of Moore's law in electronics but also thermal limits as power consumption increases with higher bandwidth and port

Integrated Photonics for Data Communication Applications. DOI: https://doi.org/10.1016/B978-0-323-91224-2.00003-5
© 2023 Elsevier Inc. All rights reserved.

counts and gradually exceeds the capability of practical IC cooling [9]. Given that future DCNs need to transmit each bit over a link at an energy of 1 pJ or less [10], it is widely accepted that new approaches are required to address these challenges.

Optical interconnect technologies with the advantages of low bit error rate, wide bandwidth, and high energy efficiency are considered as a promising way to break through the energy-bandwidth limitations of conventional electrical wires and improve interconnection performance [11,12]. At the link level, optical transmission was initially adopted for large-scale data centers between 2007 and 2010 for the transition of link data rate from 1 to 10 Gb/s [13]. The capacity of optical link has since increased to 400 Gb/s and is predicted to increase tenfold every 7 years [14]. At the switch level, optical switch offers large bandwidth switching capability. In fact, the predicted growth rate of optical switching bandwidth density even reaches tenfold every 5 years [15]. Although optical switches cannot fully replace electronic packet switches (EPS) due to the lack of optical buffering technology, it is well recognized that the combination of fast optical circuit switching and electrical packet switching is one of the most attractive solutions for the future DCN architecture [16–19]. Therefore in this chapter, we focus on the optical switching techniques that would facilitate next-generation DCNs. We start with a brief review of the key figures of merit of optical switching networks from both system and device perspectives. We then provide overviews of optical switch fabric architectures and integrated optical switching technologies. Afterwards, we discuss the technical aspects of silicon photonic-based switches, InP-based optical switches, and hybrid switches successively. Representative interconnection schemes powered by optical switching for data center and HPC networks will follow. Finally, we analyze the current technical challenges and discuss future directions.

8.2 Key figures of merit for optical switching networks

8.2.1 System metrics

Data centers consist of a multitude of servers as computing nodes and storage subsystems interconnected with the appropriate networking hardware and accompanied by highly engineered power and cooling subsystems. As introduced in Section 1.2.4, a typical data center could have 20 servers housed in each rack, 20 racks grouped into rows (clusters), and about 200 rows in total. The DCN is responsible to support the large amounts of workload exchanged among the parallel server machines. Traditional DCN normally uses a multitier architecture, such as three-tier architecture with core, aggregation, and access layers or two-tiered spine-and-leaf architecture. Different layers serve diverse purposes and employ different switching schemes. For example, the network architecture in the Facebook data centers (known as F16 fabric) utilizes three switch tiers which include top-of-rack (ToR), fabric, and spine switches from bottom to top, interconnecting all the racks within the data halls. Among them, the ToR switches interconnect the servers via copper or optical links, while the interrack communication is handled by layers of electronic switches.

Ideally, the DCN should provide full bisection bandwidth, and thus the oversubscription ratio would be 1:1 indicating high server utilization and computation efficiency. However, due to the super linear costs associated with scaling the bandwidth and port density of electronic switches, such a design would be prohibitively expensive for a large-scale DCN. In practice, DCs tend to enforce an oversubscription ration of 1:4−1:10. There is more bandwidth available for intrarack communication than interrack communication, and a similar trend would be found at higher switching layers [20].

From the system perspective, a set of stringent requirements are imposed on the switching networks, key ones are as follows.

1. Capacity: Although the traffic between users and DCs is substantial, the vast majority of the data communication is within the data center (ease−west traffic). Recent studies have shown a continuous increase in the interrack traffic with a clear majority of traffic being intracluster ($>50\%$) [21]. Higher bandwidth interconnects in combination with high-capacity switching elements are required especially for interrack and intercluster communications, to avoid the congestion drops caused by the inherent burstiness of traffic flows and intentionally oversubscribed networks [22].
2. Latency: In the closed environment of DCs, latency (defined as the time for a packet to traverse the network from sender to receiver, i.e., end-to-end latency) is dominated by the switch latency mainly contributed by the buffering, routing algorithm, and arbitration. Low latency is a crucial performance requirement especially for mission-critical and latency-sensitive applications where microseconds matter (e.g., financial networking).
3. Interconnectivity: The servers in data centers normally have 10s to 100s of concurrent active flows. A considerable number of flows will traverse an uplink at the ToR switch as interrack communication, which demands sufficient interconnectivity of the switching network to accommodate the high number of concurrent flows. Moreover, considering most flows are short and tend to be internally bursty, fast and dynamic reconfiguration of such interconnectivity (e.g., statistical multiplexing) is also needed to guarantee efficient bandwidth utilization and timely service delivery.
4. Scalability and flexibility: The network architecture should enable scaling to large number of nodes for future capacity needs in a cost-efficient manner, for example, the extension of an existing network in terms of both node count and bandwidth in an incremental fashion, that is, without having to replace a disproportionate amount of the installed hardware. On the other hand, the DCs are also expected to adopt technologies that allow them to flexibly manage the service delivery and adapt to changing needs
5. Fast reconfiguration and control mechanism: To fully benefit from the flexibility enabled by statistical multiplexing, optical switch fabrics with fast switching time together with a fast control mechanism are desired. The implementation of the switch control should follow the increase of the network scale and optical switch port-count, and more importantly, occupy as few resources as possible.
6. Power-efficiency: DCs currently consume about 1% of the entire global electricity [23]. Network power-efficiency needs to improve for reducing the energy-related costs. It is expected that for future DCNs, the energy dissipated to transmit each bit over a link will have to be brought down from several tens of pJ today to ~ 1 pJ [6].

To summarize, high capacity switching networks with low switching latency and fine switching granularity (e.g., deploying statistical multiplexing) are necessary to effectively improve the bandwidth efficiency and handle the burstiness of the DC

traffic flows. However, for current DCNs, the large number of concurrent flows makes large interconnectivity as well as fast reconfiguration a necessity for the switches, in which case the circuit-based approaches may be challenging to employ. The pairwise interconnection and tens of milliseconds reconfiguration time have also strictly confined the applications to well-scheduled and long-lived tasks. With increasing number of server nodes and rapid upgrade in *I/O* bandwidth, these requirements would be quite challenging in terms of both switching node and network architecture. First, it is difficult for the conventional electronic switch to satisfy the future bandwidth needs. Such bottleneck can be partially attributed to the imminent end to conventional Moore's Law. Along with this, despite the continuous efforts from merchant silicon providers toward the development of application-specific integrated circuits (ASICs), the implementation of high-bandwidth electronic switch node is limited by the switch ASIC *I/O* bandwidth due to the scaling issues of the ball grid array package [8,24]. Higher bandwidth is achievable by stacking several ASICs in a multitier structure, but at the expense of larger latency and higher cost. Another limiting factor is the power consumption. As an electronic switch has to store and transmit each bit of information, it dissipates energy with each bit transition, resulting in power consumption at least proportional to the bitrate of the information it carries. In addition, the O/E/O conversions and format-dependent interfaces need to be further included as front-end, greatly deteriorating the power-efficiency and cost-efficiency performance.

In addition, interconnecting thousands of ToR switches, each with a large amount of aggregated traffic, would put enormous pressure on the multitier tree-like topology employed by current DCNs. Due to the limiting performance in terms of bandwidth and port density of electronic switches, the network is commonly arranged with over-subscription. Consequently, data-intensive computations become bottlenecked especially for the communication between servers residing in different racks/clusters. The multiple layers of switches also add significant latency when a packet traverses all DCN to reach its destination, mainly caused by the queuing delay of buffer related processing. Therefore to effectively address the bandwidth, latency, scalability, and energy requirements imposed by next-generation DCNs, innovative optical switching technologies, and network architecture are of paramount significance.

8.2.2 Device metrics

Besides the system-level requirements mentioned above, there are also metrics for the optical switching devices to be considered for practical implementation in DCNs. Here, a set of key device-level metrics are discussed:

1. Switching speed: The required switching speed is very application dependent and can range from milliseconds, for backup planning applications, to nanoseconds, for certain computing applications. As a result, different optical switching technologies and architectures, such as slow millisecond switches [MEMS, beam steering based on piezoelectric, thermo-optic (T-O) switches] or fast nanosecond burst/packet switches [semiconductor optical amplifier (SOA), electro-optical based beam steering and interferometers] are needed. Nevertheless, it has been reported that reducing the switching time below the

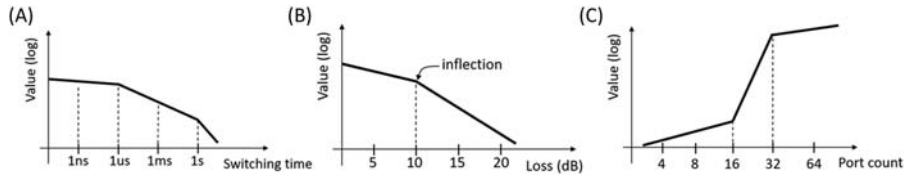

Figure 8.1 (A) Relationship between switching time and value. (B) Relationship between switch insertion loss (worst case) and value. (C) Relationship between switch radix and value.

100-ns mark does not bring significant benefit in DCNs as the retraining time for all the links traversing the switch may exceed 100 ns after the switch is reconfigured [12]. Fig. 8.1A shows the value of a switch as a function of its switching time. It can be seen that the value slowly decreases in the 1 μs to 1 s range, and significantly drops beyond 1 s.

2. Insertion loss (power penalty): The end-to-end signal loss introduced by the switching element is an important parameter that determines the system power budget, which depends on the working principle and topology of the switch. Ideally, switches should be lossless, especially when multiple switches are cascaded, otherwise in-line amplifiers might be required to compensate for the cumulative loss. Currently hybrid integration of Si/SiN photonics platforms and III−V SOAs are being investigated, representing one of the most promising solutions to accomplish lossless optical switching. Fig. 8.1B shows the value of an optical switch versus the insertion loss, indicating that a loss below 10 dB is generally desired.
3. Noise + nonlinearity: High losses require amplification stages that degrade the optical signal-to-noise ratio (OSNR) and therefore can prevent large port count switch implementation. This is particularly pronounced when SOAs are employed because they have larger noise figures which expedite the OSNR degradation. Moreover, at high input power levels, nonlinearities of the SOA can also introduce signal distortions. Careful design of SOA based switches is therefore crucial, as will be discussed later in this chapter.
4. Crosstalk: Crosstalk is another important metric for assessing the quality of optical switch. Typical crosstalk < -35 dB is required in large switch matrix. This demands for very good ON/OFF ratio optical gates in the optical switch architecture. SOA based optical gates are preferable as >45-dB ON/OFF ratio can be easily achieved for broadband operation, while microring or Mach−Zehnder interferometer (MZI) based optical gates are for narrowband operations [25].
5. Port size scalability: The scalability of optical switch is also of great importance since for applications in large-scale DCNs, a minimum number of 16 to 32 ports is normally required, while beyond this minimal requirement [10], the increase in value becomes marginal, as illustrated in Fig. 8.1C. Different switching mechanisms and architectures have drastically different scaling capabilities, for example, MEMS switches have been successfully demonstrated with hundreds of ports, while the current record for T-O MZI-based switches is 64×64.
6. Cost: To make optical switches cost-competitive, the cost associated with making a link configurable must be considerably (e.g., one order of magnitude) smaller than the cost of adding an additional link. Considering that the current price of 100 Gb/s optical links (including the cost of optical transceivers and interfaces) is continuously decreasing from the order of $1000/port to a value closer to $100 [10,26]. To keep one order of magnitude advantage, optical switches should thus target a cost per port of ∼$10.

7. Power consumption: Power consumption of the switches (e.g., for SOA, microring or MZI), is very important as it determines the number of photonic active elements that can be integrated in a single chip and therefore the scalability of the switch. In general, a voltage driven optical gate (microring or MZI) is preferable to a current driven one (SOA) because of less power consumption.
8. Control complexity: The overall control circuit complexity (including the temperature stabilization and wavelength registration in the interferometer devices) should be also carefully considered, in particular for large port count switches with 100s or 1000s of components. The complexity of the control drivers is also related to the switching speed as well as the power consumption.
9. Polarization dependence: Polarization-dependent loss (PDL), for passive devices, or polarization-dependent gain (PDG), for amplifiers, such as SOAs, should also be considered. For the short links and closed environment of the DCN, they would have similar effects as insertion loss. SOAs with low PDG ($<$0.5 dB) have been realized.

In summary, optical switches need to feature over 16 ports and $<$1 ms switching speed to be valuable. On the other hand, the tradeoff between cost/power efficiency and power penalty also requires careful consideration in the switch design. For example, silicon photonic switches can be cost-effective but normally have high insertion loss, while III−V SOA-based switches can be made lossless at the cost of higher price and power consumption. In the next sections of this chapter, more details about different optical switch architectures and technologies will be reviewed.

8.3 Overview of optical switch fabric architectures

It has been widely recognized that the switching device performance depends critically on the selection of topology, including switch blocking characteristics, crosstalk suppression, total number of switch cells and number of cascading stages [6,9]. Some of the classical switch architectures, such as crossbar, Banyan-type, Clos, and Beneš networks, are adopted from electronic switch network designs, while the others are made by pioneers in optical switch fabrics, especially those with innovations to offset the limitations of photonic integration technologies. For instance, N-stage planar architecture was proposed to eliminate waveguide crossings, path-independent loss (PILOSS) network was designed to achieve a loss uniformity across all paths, and dilated networks were exploited to completely cancel the first order crosstalk. We list common optical switch architectures, namely Banyan-type, N-stage planar, switch-and-select, PILOSS, crossbar, Beneš, and dilated Beneš (shown in Fig. 8.2A−G) in Table 8.1 based on their blocking characteristics, together with other key figures of merit, such as the level of crosstalk, total number of switching elements and number of switch stages.

8.3.1 Banyan-type blocking networks

Banyan-type switch fabric, which was proposed for use in computer networks, is also attractive for high-speed electronic and optical switching applications.

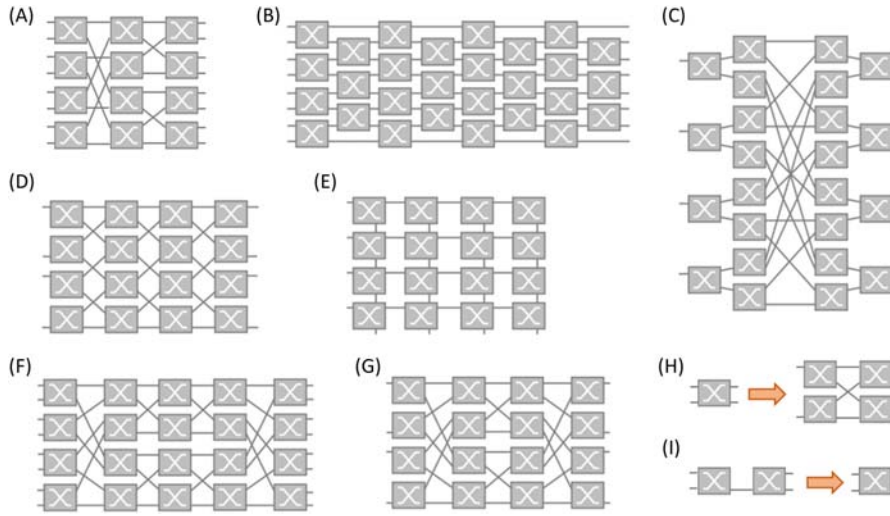

Figure 8.2 Schematic representation of optical switch topologies: (A) Banyan-type, (B) N-stage planar, (C) switch-and-select, (D) path-independent loss, (E) crossbar, (F) Beneš, (G) dilated Beneš, and (H and I) illustrations of how to substitute a 2×2 switching element with dilated four switching elements.

A Banyan-type network is defined as a class of multistage networks that have no path diversity, and there exists some variants including banyan, omega, baseline, and n-cube [27], differing in the interstage connecting pattern. The Banyan-type switch fabric provides the minimum diameter for a fully connected but blocking network, which means that any input port has a full connectivity to any output port provided that no contentions occur with existing connections in the network. It thus has the minimum number of switch cells among the switch topologies listed in Table 8.1. Banyan-type network suffers first-order crosstalk (i.e., the crosstalk from the same switching element) as any 2×2 switching element traversed by two signals at once.

8.3.2 Rearrangeable nonblocking networks

Rearrangeable nonblocking switch fabrics have so far drawn the most attention, since this type of switch networks typically require smaller number of switch cells, as revealed in Table 8.1. This is a critical consideration in photonic integration that translates into lower design and fabrication complexity. However, one should note that this type of networks can only establish connections of all permutations of input ports to output ports if rerouting existing connections are allowed.

Beneš topology is undoubtedly the most popular, because it has minimum number of switch cells to construct a nonblocking $N \times N$ network [28]. It was derived from the Clos architecture, with a bank of 2×2 switch cells forming each of the input and output stages and recursively replacing the center stage by smaller size

Table 8.1 Key figures of merit for different optical switch topologies, where N represents the port number, while X represents the crosstalk ratio of a 2×2 switching element (in dB).

Topology	Blocking characteristic	Element number	Stage number	Crosstalk in dB (worst)
Banyan	Fully connected but blocking	$(N/2)\log_2 N$	$\log_2 N$	$-X - 10\log_{10}(\log_2 N)$
N-stage planar	Rearrangeable nonblocking	$(N/2)(N-1)$	N	$-X - 10\log_{10}(N)$
Switch-and-select	Wide-sense nonblocking	$2N(N-1)$	$2\log_2 N$	$-2X - 10\log_{10}(\log_2 N)$
PILOSS	Wide-sense nonblocking	N^2	N	$-X - 10\log_{10}(N-1)$
Crossbar	Wide-sense nonblocking	N^2	1 to $2N-1$	$-X - 10\log_{10}(N-1)$
Beneš	Rearrangeable nonblocking	$(N/2)(2\log_2 N - 1)$	$2\log_2 N - 1$	$-X - 10\log_{10}(2\log_2 N - 1)$
Dilated Beneš	Rearrangeably nonblocking	$2N\log_2 N$	$2\log_2 N$	$-2X - 20\log_{10}(2\log_2 N - 1) + 3$

Clos networks until it reaches 2×2. This produces $2\log_2 N - 1$ stages with a total of $N/2(2\log_2 N - 1)$ cells. Beneš topology also encounters first-order crosstalk. Comparing to Banyan-type blocking networks, the additional switch stages not only enhance the connectivity but also provide the path diversity.

N-stage planar network is also widely chosen for building modest-scale switch fabrics [29]. As mentioned above, this design was motivated to avoid optical crossovers, as a derivative from the crossbar architecture. The N-stage planar architecture has $N/2(N-1)$ primitive 2×2 cells, and a path-dependent performance with paths traversing $N/2$ to N switch cells. The N-stage arrangement provides a strong path redistribution capability that largely increases path diversity.

The first-order crosstalk imposes stringent requirements on the design of elementary switch cells. Dilation can be used to modify the switch fabric so that each individual switch cell only carries one signal at once, dramatically reducing crosstalk (to the square of elementary switch cell), though at the expense of a greater number of switching elements and cascaded stages. The construction process of a $N \times N$ dilated Beneš network includes two steps, as indicated by Fig. 8.2H and I. First, each 2×2 switching element in a $N \times N$ Beneš switch network is substituted by a 2×2 dilated switching block formed by four 2×2 switching elements (i.e., Fig. 8.2H), creating a switch fabric with doubled stages. Then, since each 2×2 switching element only carries one signal, the even-numbered stages (except the last stage) can be combined with the following stage (i.e., Fig. 8.2I) to form one switching element, such that the stages of the finished $N \times N$ dilated Beneš switch network will be reduced to $2\log_2 N$ (i.e., one more stage comparing to the original Beneš switch network), while the number of switching elements will be $2N\log_2 N$. Considering that the fully connected cell pair can only be configured in two states, it can be seen that each 2×2 dilated block has four switching states except the last stage which can only be set in two states. Therefore its global switching states is $2^{(N/2(2\log 2N-1))}$, which is the same as a Beneš network.

8.3.3 Wide-sense and strictly nonblocking networks

Wide-sense and strictly nonblocking networks can set up paths between any idle inputs to any idle outputs without interference with existing connections. The difference lies in that routing in the former type of networks needs to follow certain rules while the latter has no restrictions. This type of networks is generally favored due to the simplified switching control system.

Crossbar and PILOSS are two typical wide-sense nonblocking networks. Both topologies require N^2 switching cells and are not fully immune to first-order crosstalk, but PILOSS wins with respect to loss uniformity since any path traverses exactly N switch cells [30]. This number, however, varies from 1 to $2N - 1$ for crossbar architecture. They both require that optical routing paths contain strictly one switch cell in bar state with the rest set to cross state, which rules out path diversity.

8.4 Overview of integrated optical switch technologies

Based on the physical effects that enable the switching process, optical switches can be classified into the following categories: electro-optic (E-O) switching, T-O switching, free-space optomechanical switching, liquid crystal-based switching and SOA-based switching. Table 8.2 shows a comparison of key performance metrics of different optical switching technologies via different platforms, such as MEMS, InP, and silicon photonics. Yet, the relative importance of these metrics should be referenced depending on the specific optical switching application, especially for optical switches that can be truly deployed in large-scale DCNs. In this section, we review the optical switch technologies implemented on mainstream integration platforms. Furthermore, based on a narrowed down 2D metric space, the development prospect of the emerging III−V-on-Si hybrid integrated switches is envisioned.

1. MEMS-based switch: As already been commercialized, MEMS-based optical switches are the most common and mature free-space switching technology. Electrostatic drivers are commonly utilized to manipulate the intensity or the direction of light beam via reflective/refractive structures or diffractive/interference effects, benefiting from energy efficiency and ease of control, though the typical voltage required is relatively high (reach up to 100−150 V) [12,31]. MEMS switch systems have been demonstrated to allow the connectivity of hundreds of ports, with low loss, low crosstalk ratio and low power consumption [32]; however, the installation and calibration with surface normal microoptics introduces considerable complexity that ultimately leads to the relatively high cost per port. Furthermore, the switching speeds for MEMS switches are on the scale of microseconds to milliseconds, which could be unsatisfactory for the low latency required by modern data centers. Their high cost and high latency make it challenging to broadly introduce MEMS switches in future DCNs.
2. InP photonic switch: InP-based integration enables complex photonic integrated circuits with high-performance active components, such as SOAs and lasers. Optical switches implemented on InP platforms have primarily relied on SOA gates in a broadcast-and-select topology, where SOA perform as either an amplifier or absorber depending on the electrical bias current, featuring the advantages of broadband performance, ease of control, and high ON/OFF extinction ratio. Fig. 8.3A. shows the first monolithically integrated 16×16 all-active SOA-based switch [33]. Fig. 8.3B presents a lossless 8×8 SOA-gated switch in a similar configuration but with a passive−active integration scheme to enhance the power efficiency and OSNR [25]. Since the SOA performance might be bounded by the chip inherent losses, one possible solution is to combine the MZIs with short, distributed SOA gates for loss compensation and crosstalk suppression. Following this strategy, a small-scale 4×4 MZI-SOA switch was demonstrated with ultrawide input power dynamic range, while the tiny on-chip loss can readily be compensated by increasing the SOA bias, as shown in Fig. 8.3C [34].
3. Silicon photonic switch: Compatible with mature CMOS manufacturing technologies, silicon photonics offers a promising platform for integrated optical switches. At present, the main research direction is interference-based switching, that is, building switch cells using MZIs and microring resonators (MRRs). For the former, a 64×64 T-O MZI-based Beneš switch has been reported, representing the current record for silicon photonic switch in terms of port count [35]. For the latter, MMRs can be reconfigured in both space and wavelength domains, hence offering resource reallocation at finer granularity. A novel

Table 8.2 Comparison table of different optical switching technologies.

	Scalability/record	**Switching time**	**Insertion loss**	**Crosstalk**	**Power consumption**	**Polarization**
Free-space 3D MEMS	Very large/ $> 1000 \times 1000$	10−20 ms	Low	Low	Medium (high driving voltage)	Independent
SiPh MEMS	Large/ 64×64	Sub-μs	Medium	Low	Medium (high driving voltage)	TE (can be independent)
SiPh T-O	Medium/ 32×32	10 s of μs	Medium to high	Medium	High	TE (can be independent)
SiPh E-O	Medium/ 32×32	∼ ns	High	High	Medium (PDE)	TE (can be independent)
III−V MZI	Small/ 8×8	∼ ns	High	High	Medium (PDE) Low (QCSE)	Independent
III−V SOA	Medium/ 16×16, $8 \times 8 \times 8\lambda$	∼ ns	Low, even lossless	Low	High[a]	TE (can be independent)
LCOS	Medium/ 1×20	100 ms	Medium	Medium	Low	TE

[a]Consumed power for both actuation and amplification.

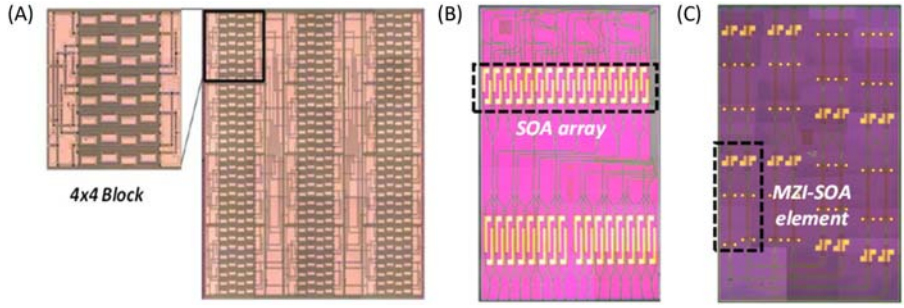

Figure 8.3 (A) InP 16 × 16 all-active semiconductor optical amplifier switch. (B) InP 8 × 8 active−passive semiconductor optical amplifier switch. (C) InP 4 × 4 Mach−Zehnder interferometer-semiconductor optical amplifier switch.

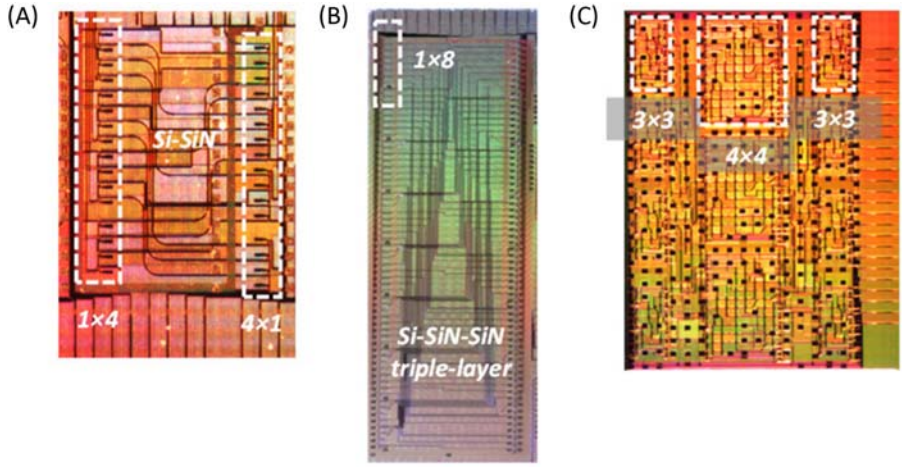

Figure 8.4 (A) Si−SiN two-layer 4 × 4 microring resonator switch. (B) Si−SiN−SiN triple-layer 8 × 8 microring resonator switch. (C) Silicon 12 × 12 microring resonator Clos switch.

switch-and-select topology based on MRRs was proposed, in which spatial multiplexers were integrated in a $1 \times N$ bus structure for signal fan-in and fan-out [36]. Based on this design scheme, a 4 × 4 switch was demonstrated, as shown in Fig. 8.4A, on a SiN-on-Si dual-layer integration platform which significantly reduced scaling overhead in the waveguide shuffle, though exhibiting over 10-dB fiber-to-fiber loss. Subsequently, a triple-layer 8 × 8 switch fabric was designed with improved insertion loss [6] (Fig. 8.4B). By using a customized Clos topology, such design can be further extended to larger port counts, and a 12 × 12 switch fabric was demonstrated, as shown by Fig. 8.4C [37]. Moreover, MMRs (overcoupled) can be used as wavelength selective phase shifters in the arms of MZIs, paving path to a new type of space-and-wavelength selective switch.

To deploy and enjoy the advantages of optical switching in next-generation DCNs, multiple key performance metrics are targeted, including fast switching

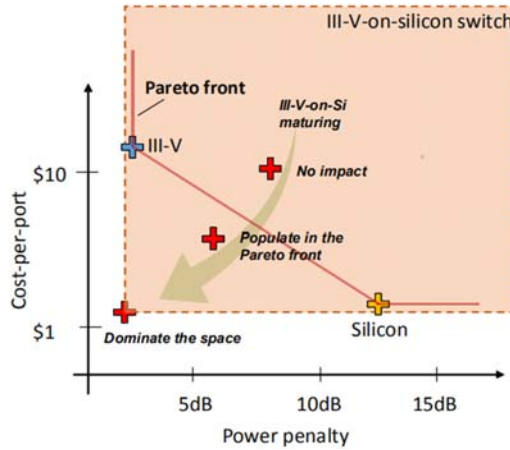

Figure 8.5 Pareto front of current III/V and silicon photonic switches versus III/V-on-Si switch.

speed, large port counts, low cost, and low power consumption. Silicon integrated switches are compact and can potentially be cost-effective, but their high insertion losses remain problematic. InP platforms permit outstanding E-O performance and energy efficiency, along with high-performance on-chip amplifiers and lasers, but the mismatching of their functional densities and wafer sizes with those in silicon CMOS production would result in high cost. As shown in Fig. 8.5, by consolidating the figures of merit to a 2D plane of cost-per-port and power penalty and plotting both InP and Silicon technologies on the plane, it can be sensibly envisioned that the combination of the strengths of InP technologies with the full capabilities of silicon photonics will transform next-generation photonic integrated switches. As the hybrid/heterogeneous integration technologies continue to mature, it is believed that III−V-on-Si optical switches will populate in the Pareto front and even become dominating. Intensive research efforts are being, and will be continuously, devoted in investigating III−V-on-silicon optical switches via various approaches, including epitaxial growth, which represents the ultimate path, die-bonding with microtransfer printing, and flip-chip bonding which have higher technology readiness levels in the near term. In the following sections, state-of-the-art demonstrations of silicon, InP, and hybrid III−V-on-silicon switches will be successively elaborated in detail.

8.5 Silicon photonic switch fabric

Continuing the discussions above, this section along with Sections 8.6 and 8.7 will introduce a variety of switch fabrics based on different material platforms, elucidating their operation principles, and highlighting their performance metrics. To start with, this section focuses on silicon photonics, which leverages the mature CMOS technologies to provide a low-cost platform for integrated optical switches. Here, the state-of-the-art T-O- and E-O-based silicon photonic switch fabrics are discussed.

8.5.1 Thermo-optic Mach–Zehnder interferometer switch fabric

8.5.1.1 32 × 32 path-independent loss switch

A 32 × 32 T-O switch has been fabricated using a 300-mm CMOS pilot line [38]. In this example, the switch matrix topology was PILOSS that consisted of an array of 2 × 2 switching elements and intersections. The 2 × 2 switch here was based on an MZI, of which the 3-dB couplers were directional couplers. Each 2 × 2 switch had a T-O phase-shifter on each arm. The T-O phase shifter was composed of a TiN heater which was located 1.5 μm above the silicon waveguide to prevent optical losses. The switching power from the bar-state to the cross-state (i.e., electrical power for π phase shift) was ≈18 mW [39]. The intersection has a numerically optimized adiabatic design [40]. The switch chip was fabricated on a 300-mm silicon-on-insulator wafer by using a 45-nm node CMOS pilot line, shown in Fig. 8.6A. The diced chip was flip-chip bonded to a ceramic interposer, in which electrode pads on the switch chip were fanned out to land-grid array. The interposer was inserted into a socket, which was assembled onto a printed circuit board with control electronics. As for the optical connection, a 5.5% silica planar–lightwave–circuit connector was attached to the switch chip. Fig. 8.6B shows the assembled 32 × 32 switch.

The average fiber-to-fiber insertion loss was 10.8 dB with a standard deviation of 0.54 dB. The bandwidth for −20-dB crosstalk was 3.5 nm and can be expanded to 14.2 nm by exchanging output ports of the switching elements [41]. The electrical power consumption of the chip was 1.9 W.

8.5.1.2 64 × 64 Benes switch

A 64 × 64 MZI-based thermo-optical switch using wide-waveguide phase-shift arms and a novel switching element calibration method [42] was demonstrated, as

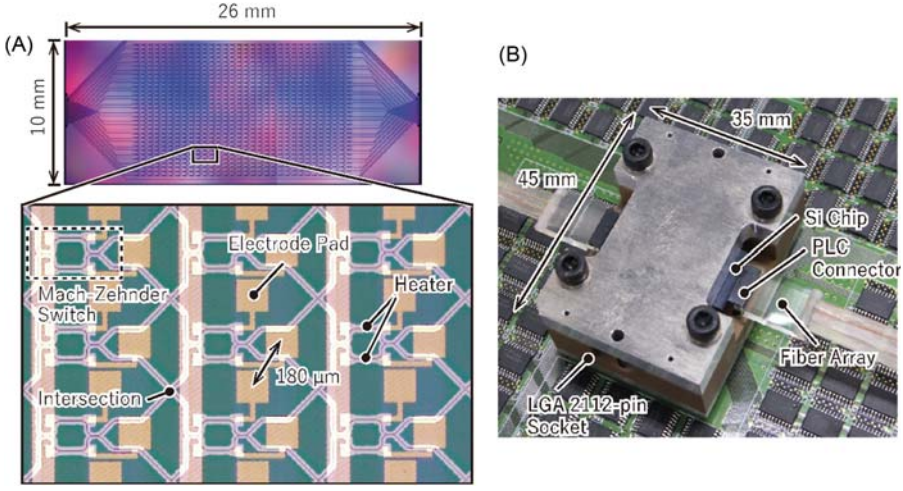

Figure 8.6 (A) Microscope image of the 32 × 32 switch chip. (B) Assembled 32 × 32 switch.

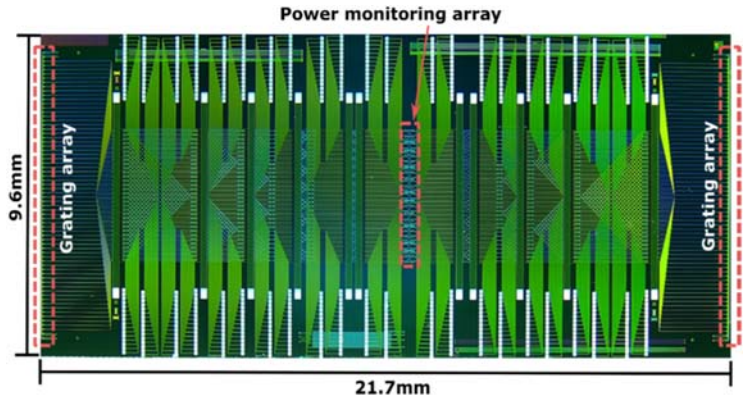

Figure 8.7 Micrograph of the 64 × 64 Mach−Zehnder interferometer-based thermo-optic switch.

shown in Fig. 8.7. Beneš network was used to construct this large-scale switch with fewer switching elements. It greatly reduced the complexity of the control circuit and made the switch easier to integrate and test.

The switching element contained two 2 × 2 MMIs and 100 μm long phase shift arms with negligible loss. to minimize network loss, the width of the two phase shift arms was 2 μm with loss of 0.215 dB/mm, while 450-nm-wide waveguides of 0.53 dB/mm were used for the low-loss 2 × 2 MMIs. Adiabatic tapers were used to connect these two widths of waveguides for low loss mode conversion [43]. Within a footprint of 21.7 × 9.6 mm, the 64 × 64 MZI-based T-O switch contained 352 thermo-optical switching elements, 1328 crossing waveguides, and 726 electrode pads. The insertion loss of the thermo-optical switch was 12.0−18.0 dB, and crosstalk was −44.3 to −30.7 dB, for one-input to one-output operating statuses over 45-nm bandwidth range [35].

8.5.1.3 Polarization diversity switch

Polarization insensitivity is one of the most important requirements for practical applications, but silicon nanowire waveguides are strongly polarization sensitive. To achieve insensitivity, there are two schemes: polarization diversity and polarization-insensitive device design. In the polarization diversity scheme, the two orthogonal polarizations are separated and propagate two switch matrices. Although two identical switch matrices are required (i.e., the footprint, number of the electrodes, and electronics for switch control are doubled), fabrication is much easier than for the scheme of the polarization-insensitive design. As an example, a polarization diversity 8 × 8 PILOSS switch was demonstrated [44], in which less than 0.4-dB PDL, less than 1.8-ps differential group delay (DGD), and less than −30-dB crosstalk were exhibited over 70 nm bandwidth. In addition, the disadvantage of the large footprint can be avoided by using the idle ports of the PILOSS topology [45,46].

For the scheme of polarization-insensitive design, all devices, and waveguides on the switch chip must have low PDL and DGD. This scheme saves space by using only a single switch matrix. However, it requires very high fabrication accuracy [47] because the effective and group indices of the high index contrast waveguide are sensitive to the structural parameters, resulting in degradation of the PDL and DGD. To ease this requirement for fabrication accuracy, one of the options is to choose a thick (≈ 1.5 μm) SOI wafer platform, which reduces the optical intensity in the waveguide core and sensitivity to fabrication-induced geometric changes. Using such a platform, an 8×8 multicast switch with approximately 0.6-dB PDL was demonstrated [48].

8.5.2 Electro-optic Mach–Zehnder interferometer switch fabric

E-O switches based on the carrier-dispersion effect are more complex than T-O switches. Injected free carriers would cause optical absorption loss and temperature increase, and because T-O and E-O effects have opposite impacts on the refractive index of a waveguide, they counteract each other and would reduce switching efficiency.

As applying high voltage to an E-O phase shifter would cause high carrier absorption loss, one solution is to set the initial state of the E-O switching element at an intermediate status (i.e., between the "Cross" and the "Bar" state), such that the required phase change would be only $\pi/2$ instead of π, reducing the requirement of driving voltage. Such approach can be achieved by decreasing the length of one output taper in the MMI structure and filling the space with a straight waveguide called an L-shifter. The optimized bias length of the L-shifter was obtained by measuring the optical extinction ratio of the two outputs of an MZI structure. In this example, it was 0.9 μm length. Power monitoring array was used to adjust the status of each switching element, the same as the 64×64 thermo-optical switch. The switch chip was packaged with four-layer wire bonding on both sides [49]. Using the above device design, a 32×32 Benes electro-optical switch array with a size of 5.2×12.1 mm was demonstrated with a 180-nm CMOS process, as shown in Fig. 8.8. All the switching elements were isolated by Si trench to reduce thermal and electrical crosstalk.

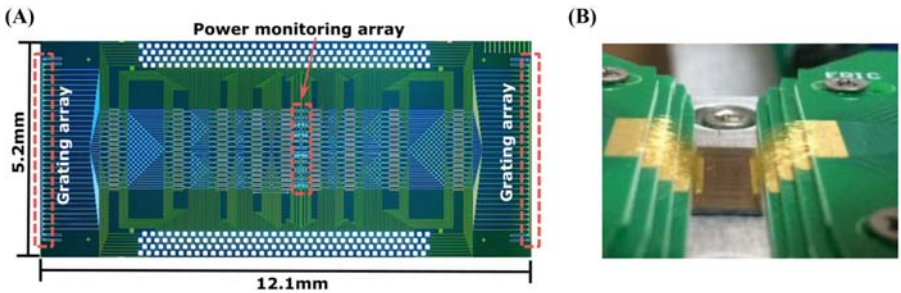

Figure 8.8 (A) Micrograph of 32×32 Mach–Zehnder interferometer-based electric-optic switch with improvement in the crosstalk by a phase bias. (B) Packaged switch chip with four-layer wire bonding.

The on-chip loss, in the "all cross" status and "all bar" status, were 12.9−16.5 dB and 14.4−18.5 dB, with crosstalk of −17.9 to −24.8 dB and −15.1 to −19.0 dB, respectively, at a wavelength of 1520. The total power consumption of the device in the "all cross" and "all bar" statuses were 247.4 and 542.3 mW, respectively. Falling and rising times of 1.0 and 1.2 ns of the switching element with a phase bias were also demonstrated [50].

8.6 InP optical switch fabric

InP optical switch fabrics based on SOAs which are used for both gating and amplification operations feature low loss, broad bandwidth, high extinction ration, and fast switching speed. The opportunity to cointegrate both gain and passive elements on chip make them suitable for compact and large-scale integration [51].

8.6.1 All-active 16 × 16 semiconductor optical amplifier space switch

Multistage architectures involving cascaded SOA elements certainly enable larger switch fabrics but require additional examination of architecture design. A hybrid multistage architecture implementing broadcast-and-select substages was proposed [52] to carefully establish a tradeoff between the number of stages and the losses incurred within each switch stage. The first monolithically integrated three-stage 16 × 16 SOA-based switch with 480 integrated components was subsequently demonstrated by the same group (Fig. 8.4A) [33]. This regrowth-free all-active device applied compact total internal reflection mirrors and leveraged amplifying waveguides to provide lossless operations.

8.6.2 Active−passive 16 × 16 semiconductor optical amplifier space switch

The same port-count of 16 × 16 input/output ports was then demonstrated on an active−passive cointegration platform [53] with the intent to allow further enhancements in terms of optical power handling and noise performance. This switch is designed implementing a hybrid Beneš-Tree architecture, where the central five stages are collapsed to one stage by using a 4 × 4 broadcast and select switch architecture. The input and output switching stages are implemented with 2 × 2 broadcast and select switching elements. The first and the last stage of SOAs are used as preamplifiers and booster amplifiers, respectively. The center stage enables path fast reconfigurability and is replicated in two planes to allow rearrangeable nonblocking interconnection. Fig. 8.9 shows a photograph of the switch circuit. The complete circuit included a record total number of 480 components [54] within an area of 4.0 × 13.2 mm. The three cascaded SOA gate stages are implemented across Fig. 8.9: Waveguides are visible as black lines against the green background. Gold

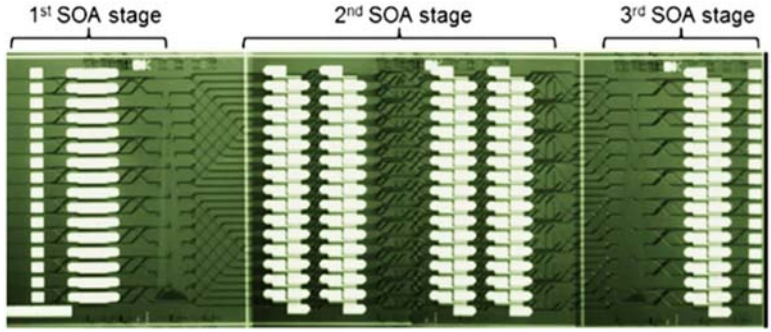

Figure 8.9 A composite photograph of the active−passive InP 16 × 16 switch.

electrodes for the 192 SOA gates are visible as six columns of 30 μm wide and 500 μm long cointegrated quantum-well active islands on a selective area epitaxial regrown InP/InGaAsP wafer.

The initial characterization is performed with a control interface to reconfigure the routing and power equalization between multiple connections to study the component level and path level operation. The shortest and the longest paths of the chip are tested for full loss map understanding. The use of low-loss passive waveguides in combination with high-contrast active gates leads to promising system level metrics in terms of OSNR with a best-case value of 28 dB/0.1 nm. However, due to the moderate scale of integration, the space switch scheme presents high path losses: The fiber-to-fiber signal losses increase as path length and complexity increase, because of the additional traversed components. Specifically, the measured total passive losses along the 32 paths include an inherent 30-dB loss mainly due to the broadcast-select capability and an excess 37-dB loss resulting from component-level imperfections, with a maximum loss variation of 12.8 dB. The gain from the SOA gates can compensate much of the excess component losses. The circuit-level loss and noise analysis are conducted by evaluating the losses in comparable circuits and through direct measurement, assuming for each of the three stages of the SOA gates in the circuit a mean gain of 12.9 dB and including a mean amplified spontaneous emission power density of about 47 dBm/0.1 nm, contributed at each SOA gate. Both the input and central SOAs were operated far from the optical saturation power. The agreement between measurements and component loss analysis is excellent, giving confidence that the loss and noise performance of the circuits can be clearly attributed to the excess losses of the components. Compared with an all-active approach, the active−passive integration scheme enables order-of-magnitude noise performance improvement as well as reduced energy requirement.

8.6.3 Active−passive 8 × 8 × 8λ semiconductor optical amplifier space and wavelength switch

Wavelength routing switching is an alternative technology to implement photonic switching architectures. Orders of magnitude increases in reconfigurability become

possible when the use of wavelength division multiplexing (WDM) is combined with on-chip wavelength selective routing. Notable examples include the monolithic tunable optical router (MOTOR) device that is an 8×8 monolithically integrated InP optical switch based on tunable wavelength converters [55]. The data center optical switch project is also employing tunable wavelength converters and arrayed waveguide grating router [56]. The integrated router interconnected spectrally project demonstrated a multistage architecture based on a combination of space switches and wavelength routing switches that utilizes fast wavelength converters and 40×40 arrayed waveguide gratings (AWG) [57]. These architectures are limited by the quality of the wavelength converted signals, the restricted dynamic range operation, and high power consumption despite the circuit compactness and high scalability. Combining space and wavelength switching circuits, photonic switches can ultimately implement circuits with wavelength, space, and time switching operation.

Stabile et al. have demonstrated an elegant example of an in-plane circuit realized on InP. It is an 8×8 integrated cross-connect on chip, which provides a connectivity of up to 64 input and output channels [58]. The architecture for broadband photonic and wavelength selective cross-connection is shown in Fig. 8.10A. The circuit consists of a combination of a broadband photonic switch stage (PSS) for the input port selection and a wavelength selecting stage (WSS) for color selection, both based on the use of a total number of 136 SOAs. Eight WDM input ports

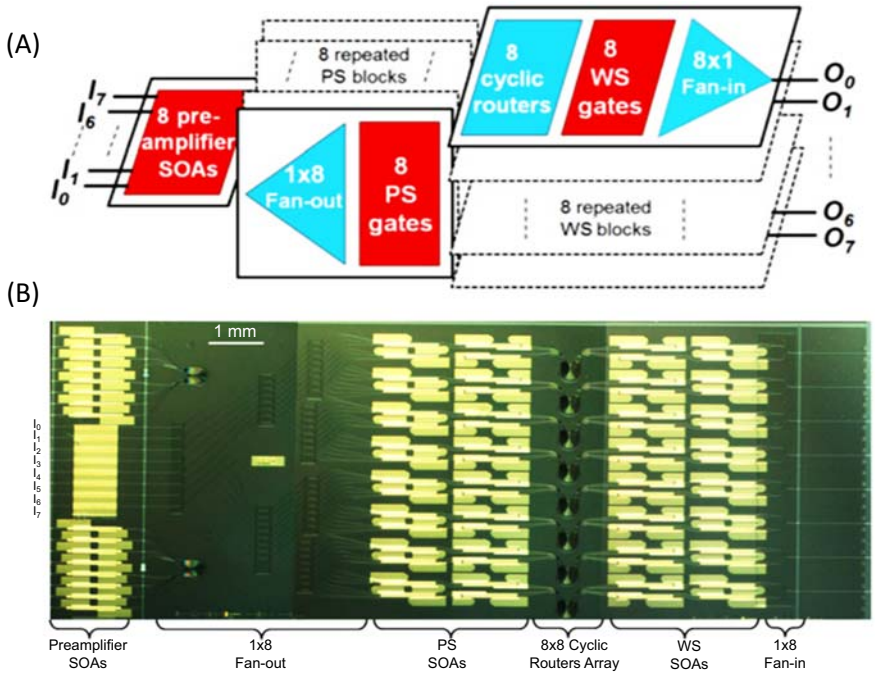

Figure 8.10 (A) Architecture and (B) photograph of the fabricated 8×8 cross-connect.

connect to an array of preamplifying SOAs and traverse a broadcast shuffle which maps each of the input to the eight output ports. The PSS SOA gates connect the inputs to eight cyclic AWGs. The AWGs route the demultiplexed wavelength channels to the second stage array of broadband SOAs to perform wavelength selection. Selected channels are aggregated with broadband fan-ins based on cascaded 1:2 MMI splitters to provide the eight colorless outputs. The gray-scale layers in Fig. 8.10A represent the other seven identical layers implemented in the integrated circuit for both stages. The red boxes are the SOA gates used in each of the stages. The device is realized on a regrown active−passive InGaAsP/InP epitaxy. An optical image of the chip is shown in Fig. 8.10B. The total footprint of the switch is 14.6×6.7 mm^2. 84% of the paths from the input to the output side of the chip are electrically and therefore optically connected in this first prototype.

Both space domain switching and wavelength domain channel selection within the same monolithic circuit are performed to show the highest level of connectivity and data capacity on a single InP photonic integrated chip [59]. OSNRs greater than 27.0 dB for 0.1-nm resolution bandwidth and a mean 13.3-dB on-chip loss are measured for each of the channels of a multiplexed channel, as evaluated at a representative path from input 0 to output 0. Multipath WDM 10-Gb/s data routing is also demonstrated. To test the simultaneous WDM multipath data routing, the wavelength multiplexed data is split into four copies using broadband splitters and launched into four input ports. Sixteen channels are routed within the 8×8 cross-connect towards output port 0. All four input paths are sequentially enabled with a round-robin scheduling and loaded with WDM input signals for dynamic multipath WDM reconfigurability studies. Fig. 8.11 shows the time traces of the output WDM

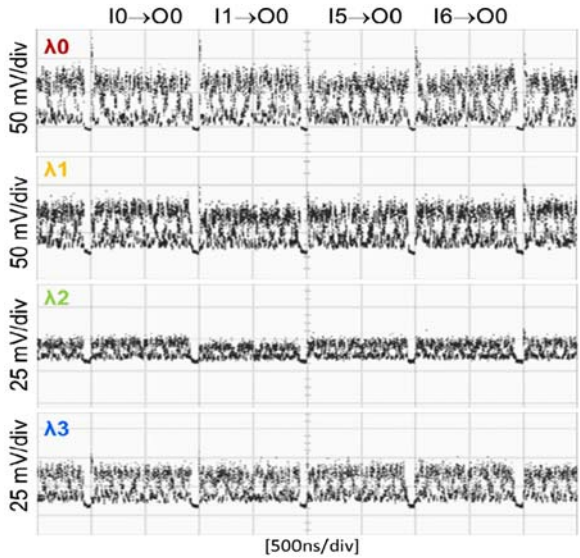

Figure 8.11 Time traces for all the wavelength of the four wavelength division multiplexing input signals.

signals. Along the *x*-axis the selected input port changes at each time slot as for the round-robin scheduling. The four sets of wavelengths are selected using external optical filters and are displayed as four separate graphs, showing well-resolved time traces. The switching occurs within 5 ns [59].

The first investigation on the impact of varied switch loads from unicast to multicast to broadcast on device performance was also performed for this chip [60,61]. Electrical power dissipation may be estimated from the mean bias current (50 mA) and the mean voltage across the SOAs (1.2 V). The preamplifiers and wavelength-select stage SOAs are operated continuously, and the space-select stages are scheduled. This leads to an electrical input power of 0.78 W for unicast and 2.88 W for broadcast operation. If dynamic scheduling was implemented for the wavelength-select stages as well, a more modest 1.44 W electrical power would be required with only 1 dB gain reduction. Further energy savings are also feasible through optical circuit level optimizations [59].

8.7 Heterogeneous photonic-integrated switch fabric

Heterogeneous photonic integrated switch fabrics enjoy the strengths of both silicon and InP photonic technologies, offering great potential for next-generation lossless switch with low cost, high switching speed, and large port court. In this section, the mainstream heterogeneous integrated switches based on flip-chip bonding approach are introduced, along with the implementation of silicon−SOA hybrid wavelength selective switch (H-WSS) fabrics. Besides, the idea of combining silicon photonics with nonvolatile materials, such as, $Ge_2Sb_2Te_5$ (GST), to achieve integrated nonvolatile switches is also discussed.

8.7.1 Gain integration with flip-chip bonding

Silicon optical switches are passive devices, so that the larger the port count, the higher the loss, from which the multiport switches suffer. A straightforward approach for loss compensation is to integrate a gain device, such as SOAs with silicon photonics. In [62], the SOI and BOX layers are etched to form a terrace, on which an InP-SOA can sit, as shown by Fig. 8.12. Then, the SOA is bonded with AuSn solder to the terrace in an upside-down ("flip-chip") orientation. For low-loss coupling, spatial mode of the silicon waveguide is expanded with a spot-size converter to match that of the SOA. For a 1×8 passive splitter integrated with the SOA, the observed fiber-to-fiber loss is 4.6 dB. Its breakdown is coupling loss of 4.6 dB (2.3 dB/interface $\times$ 2 interfaces), loss of 9 dB by power split at three stages, waveguide loss of 1 dB, and SOA gain of 10 dB. Later, by further optimization of the integration process, loss in fiber-to-fiber transmission turns to gain in a 4×4 switch integrated with a 4ch SOA [63]. The WDM transmission characteristics of this device are evaluated in [64]. Recently an 8×8 switch was integrated with an 8ch SOA [65].

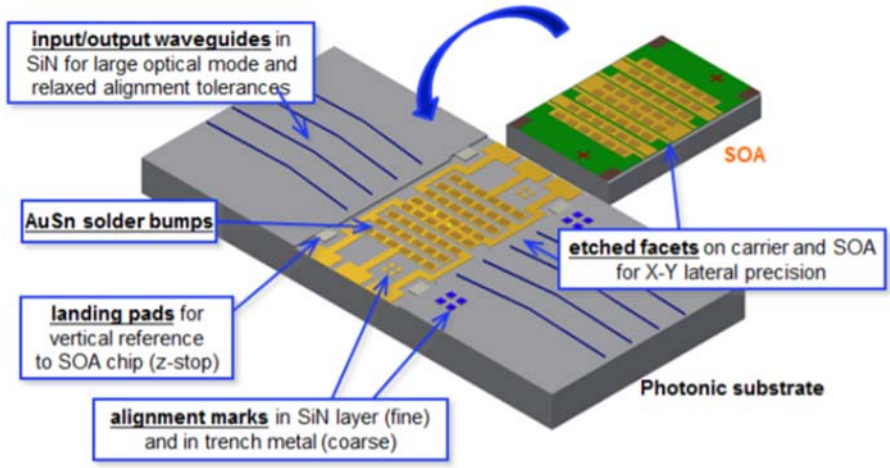
input/output waveguides in SiN for large optical mode and relaxed alignment tolerances
SOA
AuSn solder bumps
etched facets on carrier and SOA for X-Y lateral precision
landing pads for vertical reference to SOA chip (z-stop)
Photonic substrate
alignment marks in SiN layer (fine) and in trench metal (coarse)

Figure 8.12 Semiconductor optical amplifier/photonic substrate integration test vehicle (via flip-chip bonding).

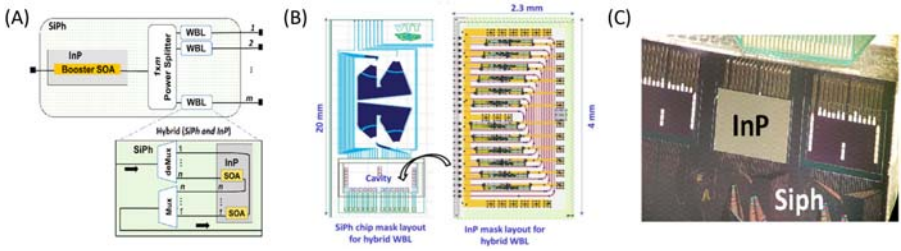

Figure 8.13 (A) Schematic representation of modular hybrid integrated wavelength selective switch. (B) Mask layout of SiPh passives and zoomed out InP actives. (C) Photograph of hybrid integrated wavelength blocker.

8.7.2 Silicon-semiconductor optical amplifier gate wavelength selective switch

The implementation of H-WSS involves silicon photonic deMux/Mux AWGs and InP SOA switching gates which are then cointegrated via flip-chip bonding. Fig. 8.13A shows the schematic of a modularly integrated H-WSS. It is based on broadcast-and-select scheme, in which the WDM signal is broadcast by a $1 \times m$ splitter and is selected by m wavelength blocker (WBL) modules at the output ports. A hybrid WBL with n-channels uses $1 \times n$ deMux/Mux silicon photonic AWGs and n SOA switching gates for the n wavelength channels. The SOA switching gates are used to pass/block the desired wavelength at each of the output ports. The port count m of the WSS is scalable by adding a new WBL switching module. A hybrid WBL module of $n = 12$ wavelength channels (10 functional and 2 test

channels) with a channel spacing of 0.8 nm (100 GHz) is implemented. The mask layouts of 1×12 AWGs on SiPh and an array of 12 SOA gates on InP are shown in Fig. 8.13B. The SOA switching gates on InP are placed in a U-shape where the input/output straight waveguides are accessed from the same side. The silicon waveguide is based on a 3×3 μm VTT's waveguide technology and provides polarization insensitive performance [66]. The InP waveguide is based on a 0.5-μm-thick waveguide InGaAsP core. Because of the size difference of the two waveguide cores, a tight alignment accuracy of less than ± 0.5 μm in the vertical and lateral directions and an air-gap less than 2 μm between the two waveguides is required to keep hybrid coupling losses below 3 dBs. In the FCB assembly process, the InP chip is directly bonded with the SiPh cavity via a thermal compression technique. More details on the H-WSS and its flip-chip bonding assembly process are reported in [67]. Fig. 8.13C shows the photo of the hybrid WBL switch after assembly.

Experimental characterizations show that the 3-dB bandwidth of deMux and Mux silicon photonic AWGs of 50 GHz and matches the designed central wavelength channel. Typical measured ON/OFF ratio is around 34 dB. Fig. 8.14 shows the gain characteristics of the gate SOA for Ch3 for varying current (mA) and varying input power. The measured SOA gain in the ON state is around 16 dB at 100 mA. BER measurements at 10- and 20-Gbps NRZ data reported error-free operation with power penalties ranging from 0.3 to 2.5 dB for different channels. These successful data transmissions through the modularly designed hybrid WBL, along with the fast nanosecond switching and amplification provided by the SOA based optical gates, verify its feasibility to be used in a low-latency optical DCN architecture, such as the OPSquare architecture, previously discussed.

8.7.3 $Ge_2Sb_2Te_5$-on-silicon nonvolatile switch

The idea of nonvolatile operation of the optical switches is attractive. It is natural to seek to combine silicon photonics with nonvolatile materials. A chalcogenide in

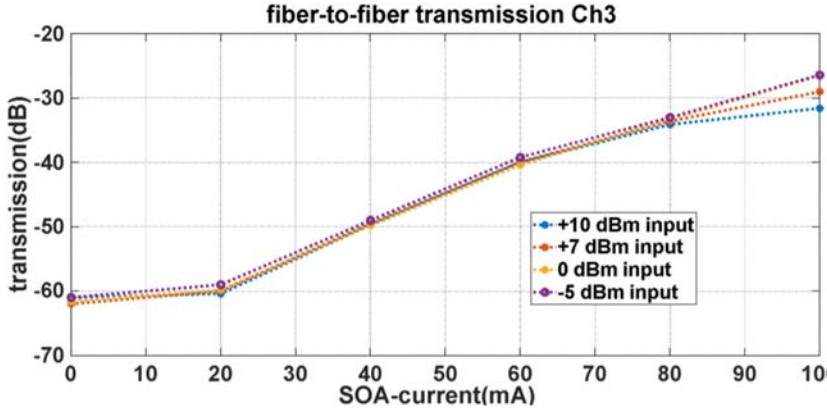

Figure 8.14 Measured fiber-to-fiber transmission and gain characteristics of semiconductor optical amplifier (Ch3).

the form of $Ge_2Sb_2Te_5$, which we denote as GST for brevity, is one example. GST has been used as optically recording medium of CD (and DVD)-RW. GST has two stable states, which differ in material index. For GST-loaded silicon waveguide or 1×1 MMI [68], it has been observed that transmission varies between two levels depending on the state of GST, making it an optical gate switch. The phase-change occurs under a proper thermal condition, controlled by laser irradiation [68] or by an integrated micro resistive heater [69]. Another example is a VO_2-loaded silicon waveguide [70]. Nevertheless, in both GST and VO_2, the imaginary part of the refractive index is large, making them unsuitable for applications to phase-shifter. Alternatively, another novel combination of silicon photonics integrated with a magneto-optical material has been recently demonstrated to realize nonvolatile optical phase shifter [71].

8.8 Data center/high-performance computing optical interconnect architectures based on optical switches

Following the previous discussions regarding different optical switching technologies, in this section, we list some representative demonstrations and testbed implementations of optical switches in the DC/HPC interconnection networks.

8.8.1 Distributed deep learning training based on optical switching

The rapid growth of data-massive applications, such as machine learning and deep learning (DL), are consuming an increasing proportion of the computation in today's data center/HPC networks [72]. The typical approach to speed up the training process of these larger DL models is parallelization using many GPU-equipped nodes, which requires high-bandwidth interconnects to support the communication requirements between training devices. As the DL models and datasets grow, distributed deep learning (DDL) is being adopted. In DDL, training models or datasets are partitioned and distributed onto multiple GPUs or TPUs. Applications often receive an allocation of resources on a set of distant and noncontiguous nodes (system fragmentation). Such trends have shifted the performance bottleneck from the compute to the network interconnect. This places a tremendous challenge on interconnect designs to provide high-bandwidth and low latency networking to sustain the continual growth of these hardware-driven DL applications (see Chapter 1 of this monograph for more details).

To address such challenge, researchers have been motivated to explore flexible network architectures based on optical circuit switches (OCSs), which offer a promising approach to reconfigure the interconnects. An example of this approach is shown in [73] where the authors use the OCS to (1) regroup a set of distant and noncontiguous nodes and (2) steer bandwidth at higher layers for efficiency. Depending on underlying traffic patterns of the nodes at different times, optimized

topology connections can thereby be dynamically formed on demand. A system architecture is proposed that leverages silicon photonic switch-enabled server regrouping using bandwidth steering to tackle the challenges and accelerate distributed DL training [73]. By using silicon photonic OCSs between servers and ToR electrical packet switches (EPSs), this architecture allows servers with intense communication requirements to be grouped locally within the same ToR switch, thereby reintroducing traffic locality between physically distant servers. The network control plane consists of (1) a Ryu-based software-defined networking (SDN) controller that manages the flow tables on the EPSs and (2) a TCP/IP client program that sends new reconfiguration requests to the SiP OCS subsystem. To demonstrate the performance of running distributed DL workloads in the proposed architecture, a 16-node HPC/data center testbed with a SiP switch-enabled reconfigurable fat tree topology is built [74], as shown in Fig. 8.15A. 1×2 and 1×4 MRR-based OCSs are used to perform server regrouping and bandwidth steering above the ToR EPSs, as shown in Fig. 8.15B. Fig. 8.15C shows the overall photograph of the EPSs, CPU servers, GPU servers, silicon photonic switches, network controller, and switch controller. Comparing to static nonreconfigurable fat tree, the switch-enabled reconfigurable architecture is experimentally measured to feature a $3.6\times$ improvements regarding the network performance. Further large-scale simulation

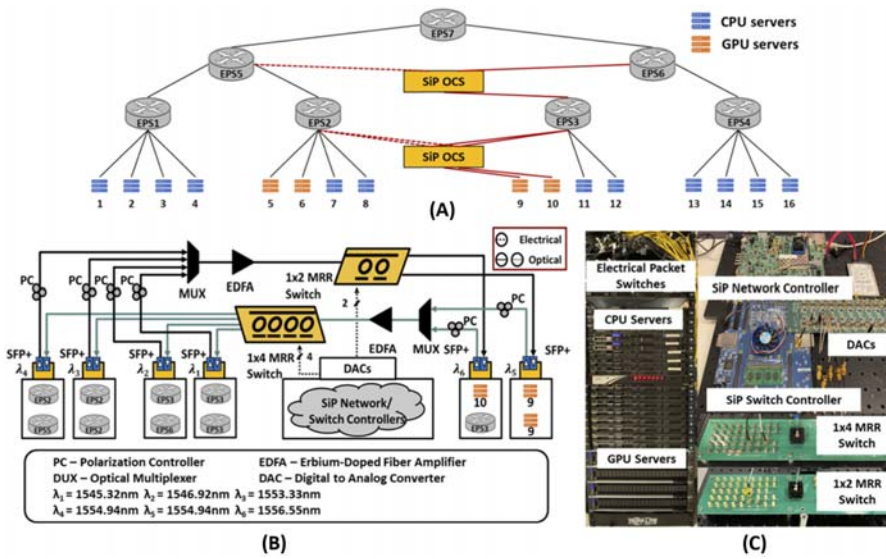

Figure 8.15 (A) A 16-node experimental testbed with SiP optical circuit switches and electrical packet switches in a reconfigurable fat tree topology. (B) Experimental setup demonstrating the cases of server regrouping and bandwidth steering above the top of rack. (C) A photograph of the electrical packet switches, CPU servers, GPU servers, SiP switches, SiP network controller, and SiP switch controller [73].
Source: From Z. Zhu et al., Distributed deep learning training using silicon photonic switched architectures, APL Photonics 7 (3) (2022) 030901.

results show that server regrouping can deliver up to 2.3× flow throughput improvement for a 2× tapered fat tree and a further 11% improvement when higher-layer bandwidth steering is employed. These results prove the feasibility of integrating optical switches into data centers and HPC systems to accelerate distributed DL training.

8.8.2 Improved resource utilization through bandwidth steering enabled by optical switching

In the supercomputer, where torus topologies previously dominated, high-radix packet switches have made hierarchical topologies, such as the dragonfly and the fat-tree more popular. The Dragonfly topology provides high connectivity with all-to-all global links at the intergroup level and aims at minimizing the number of long distance links to reduce cost [75]. However, the advantages of high connectivity are diluted by low per-link bandwidth. The bandwidth of intergroup (global) links, carrying the traffic between Dragonfly groups, can become the bottleneck for an entire network.

A major reason for this bandwidth bottleneck is due to the highly skewed traffic characteristics of HPC applications. These traffic patterns concentrate traffic on only a small percentage of links, so that only a few links are congested while most others are severely underutilized. Thus the current, best-for-all approaches using static, overprovisioned networks have topologies that are mismatched with the applications that operate over them and will likely become a bottleneck for next-generation Exaflop platforms [76,77].

Designing networks that properly balance the traffic is challenging as overprovisioning the network incurs unnecessary cost and energy, while underprovisioning leads to limitations on system performance due to the data-starved processors. In this section, a study on Flexfly is presented, in which a photonic architecture is adopted to trade global links among dragonfly groups via low-radix silicon photonic switches, allowing the network topology to be dynamically reconfigured to match the HPC application traffic [78]. In other words, the optical switch is used to steer bandwidth where it is most needed. An experimental 32 node Flexfly prototype was built with four racks connected through a silicon photonic switch. The interconnect reconfiguration time was 820 ns. The network architecture consisting of the control and data planes is shown in Fig. 8.16. The control plane is a Ryu-based SDN controller, and each rack sends a unique wavelength through the EPS. By tuning the rings in different ways, the device allows different servers to be connected bidirectionally and thus act as a wavelength and spatial OCS. When traffic is added between servers, the control plane detects the changes and initiates the configuration updates. The configurations provide a direct connection for traffic between particular servers, so that they are able to reach near full link capacity, while the other flows must compete with the background traffic and are therefore limited in throughput. The results demonstrate that the concept of bandwidth steering is a promising solution for improved network efficiency and resource allocation.

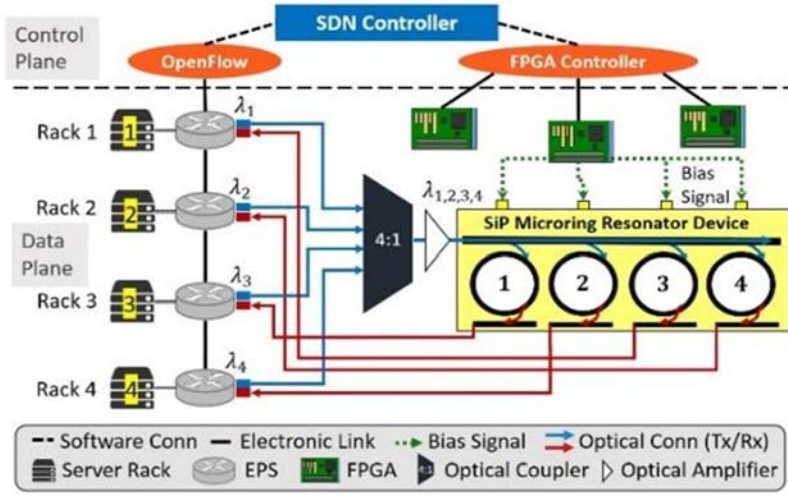

Figure 8.16 Network architecture showing the connections of the servers and top-of-rack electrical packet switches to the SiP microring resonator device [78].
Source:From Y. Shen, M.S. Glick, K. Bergman, Silicon photonic-enabled bandwidth steering for resource-efficient high performance computing, 10946 (2019) 17−25.

8.8.3 Disaggregated architecture based on nanoseconds optical switching

Recently emerging cloud applications proliferate and are becoming increasingly differentiated, imposing diverse resource requirements to DCNs. A scalable and low latency server-centric optical DCN architecture (OPSquare) based on distributed nanoseconds optical switches and control was demonstrated in [79]. However, in current server-centric DCNs, all the hardware resources (CPU, memory, accelerator, and storage) are fixed in the server-unit, which results in the "performance wall" and resource underutilization. For memory-intensive applications (like Memcached memory-caching system), there may not be enough available memory in the single server-unit to deploy the entire system, but swapping memory data out to the slower storage also results in a degradation in application performance. On the other side, the servers running accelerator intensive applications (like machine learning training) may not have enough accelerator resources to accelerate the training, while accelerator resources of servers running memory-intensive applications are unutilized. Moreover, these underutilized hardware resources also lead to a waste of capital cost and power consumption. Aiming at solving these issues of current server-centric DCNs, novel disaggregated architectures has been proposed for next generation DCNs, decoupling these hardware resources from the server unit. Based on resource pools of CPU, memory, accelerator, and storage, the disaggregated architecture guarantees flexible resource provisioning according to application requirements and improves power efficiency. The challenge is the implementation of low latency and high bandwidth scalable interconnect network and its dynamic

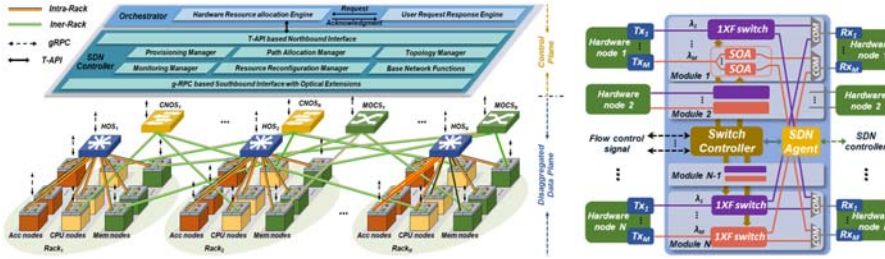

Figure 8.17 The architecture of the low latency disaggregated data center network architecture (DACON) based on distributed nanoseconds semiconductor optical amplifier-based optical switches.

and real-time automated orchestration control system to manage the resource provision and reallocation. An experimental low latency rack-scale disaggregated DCN architecture (DACON) based on distributed nanoseconds SOA-based optical switches (NOS) and nanoseconds control system empowered by an SDN control plane was proposed in [80]. The architecture is shown in Fig. 8.17. All the computing hardware resources (including CPU, memory, and accelerator) are decoupled as separated nodes and connected by different optical switches based on corresponding traffic requirements. The latency sensitive traffic between CPUs and memories is achieved by the nanoseconds optical switches, while the throughput sensitive CPU/memory-accelerator traffic utilizes slower optical switches. Using the customized SDN orchestration plane, DACON achieves flexible resource provisioning and dynamic resource reallocation. Experimental results demonstrate that the traffic processed by the DACON prototype has a minimal node-to-node latency of 378.6 ns and up to 1.74× faster application runtime compared with the server-centric architecture. Based on the computing resource reconfiguration with customized SDN orchestrator, the node-to-node latency is reduced by 21%. Numerical scalability analysis indicates that DACON achieves a CPU-memory latency of 404.8 ns for intrarack traffic and 521.9 ns for interrack traffic in a 6144-node network and memory access ratio of 0.9. Moreover, the cost and power consumption comparison shows that DACON consumes up to 34.1% less power than server-centric architectures.

8.8.4 Highly scalable fast optical circuit switch networks

In the above, the discussions mainly focused on the system architecture. In this section, the implementation of highly scalable fast OCS networks is the focus. As optical switches are incapable of offering full packet-switching functions, they always require an additional control plane to monitor the exact topology status of the physical layer, that is, the functions and connections of optical switches as well as their states of signal integrity. For large-scale fast optical switches, it is essential to develop a machine-readable mathematical modeling to generally describe the state of the physical OCS networks, such that the control operation can be automated [7].

Apart from the control plane aspects, it is also challenging to realize highly scalable OCS networks in terms of signal integrity. One straightforward way for creating scalable and strictly nonblocking optical switch networks is the multistage connection of optical switches, for example, Clos network. Considering a three-stage Clos network as an example, as shown in Fig. 8.18. The Clos network is configured by connecting r $m \times n$ switches at the first stage, m $r \times r$ switches at the second stage, and r $m \times n$ switches at the third stage. The scalability can be further expanded by recursively connecting the three-stage Clos network. For instance, strictly nonblocking 512×512 and $131{,}072 \times 131{,}072$ Clos network can be realized by three and nine-stage connections of 32×32 switches, respectively [81].

A recent work has demonstrated the capability of nine-stage cascaded operations of 32×32 silicon photonics switch for 16-channel $\times$ 32-Gbaud quadrature phase shift keying (QPSK) signals [82]. Fig. 8.19A illustrates the experimental setup. The transmitter generated 16-channel of 32-Gbaud single-polarization QPSK WDM signals spaced at 37.5 GHz from 1545.023 nm (ch.1) to 1549.515 nm (ch.16). The resulting net rates were 0.95 Tb/s (2 bit/symbol $\times$ 0.93 $\times$ 16ch. $\times$ 32Gbaud). The WDM signals were launched into the device under test (DUT) block, as shown in Fig. 8.19B. The output signal was detected by a digital-coherent receiver and demodulated to analyze the quality. To emulate multistage operation using one switch module, signal transmission was done with a circulating connection. Fig. 8.20 shows the measured Q-factors of QPSK signals at each stage, where the channels of #1, #8, and #16 were tested. Due to the accumulation of ASE noise and crosstalk, the Q-factor gradually degraded in accordance with an increase in the number of stages. The Q-factors of all the measured channels exceeded the Q-limit of 8.5 dB, assuming a 7% overhead HD FEC threshold. This shows that a 125-Pb/s

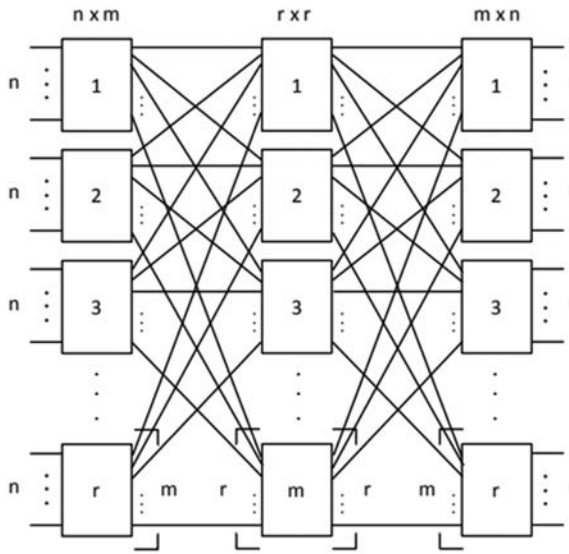

Figure 8.18 A three-stage Clos network architecture.

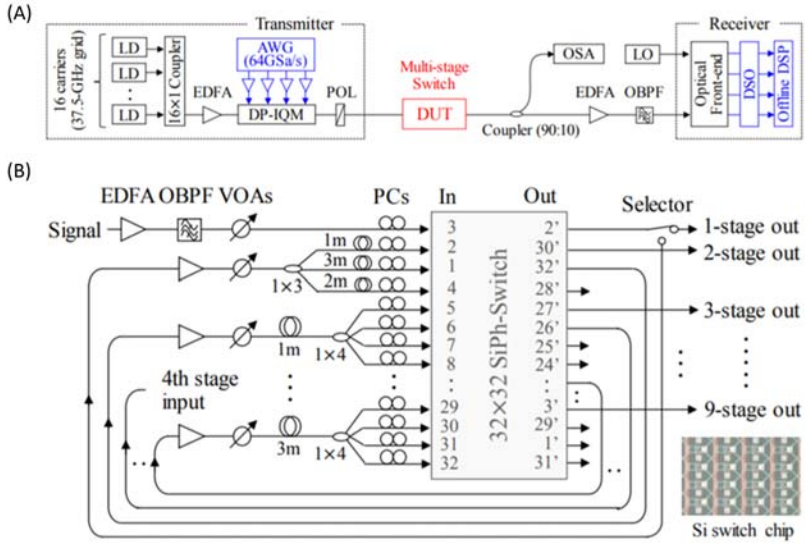

Figure 8.19 (A) Experimental setup of the nine-stage cascaded 32 × 32 silicon photonics switch. (B) The configuration of device under test block.

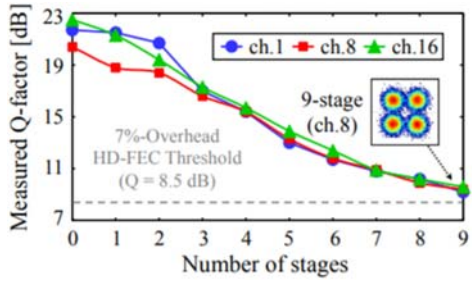

Figure 8.20 Measured Q-factor of quadrature phase shift keying signals.

(131,072 ports × 0.952-Tb/s QPSK) strictly nonblocking optical switch can be built over a nine-stage Clos network based on silicon photonics switches.

8.9 Remaining challenges and future perspectives

In this section, the underlying obstacles for the deployment of optical switching technologies in future DCNs are discussed.

First, there are many interdependent aspects in the packaging of optical switches, such as thermal, electrical, optical, and mechanical. Moreover, calibration and testing are also critically important. In terms of thermal, temperature change may have

significant impact on wavelength-sensitive switches based on ring resonators and AWGs. As for electrical design, one of the major obstacles to scaling switch port count is electrical trace fan-out that becomes increasingly complex as the number of switching elements or radix of each switch increases. In terms of optical, the technology for attaching arrays of more than 32 fibers has been established at a mass production level with respect to arrayed waveguide grating devices, etc. As for mechanical design, mechanical integrity against thermal shock and vibration has yet to be fully assessed to achieve the required reliability. For switch calibration and testing, a main concern is throughput, that is, the time it takes to complete these processes. In fact, there are N^2 combinations for the switch paths in a $N \times N$ matrix switches. The measurement time for each path and switching speed directly limit throughput. To increase throughput, test time reduction approaches, such as test sampling and lower resolution/higher speed measurement methods, that still provide sufficient data accuracy would be required.

Second, the main optical characteristics that limit the scaling of an optical switch system are the fiber-to-fiber insertion loss and crosstalk. As for fiber-to-fiber insertion loss, the largest contribution is the coupling loss between the optical fiber and the spot-size converter on the chip due to mode field profile mismatch, a more detailed discussion on fiber assembly can be found in Chapter 12. The use of 5.5% silica planar–lightwave–circuit connector would be one solution since its mode field diameter is matched to those of the silicon edge couplers and high-index fibers fusion-spliced to standard single mode fibers [83,84]. To reduce optical crosstalk, it is important to suppress leakage from other routes on the established path. To achieve this, some switch topologies (e.g., switch-and-select [85], double-layer, double Mach–Zehnder PILOSS) have been proposed. However, since the number of switching elements generally increases to improve crosstalk, it is necessary to carefully choose the switch topology that balances acceptable performance and size.

Third, the control power consumption is an essential consideration in optical switch design. For example, in switch topologies using MZI unit switches, most of the power consumption on the chip comes from the compensation of the initial phase error of the MZI switches. This is because the number of switches that are subject to π phase shift (in the "cross" state) is much smaller than the number of switches that are operating at the zero-bias point (in the "bar" state). Adjustment of a MZI switch to operate exactly at the zero-bias point is necessary to suppress crosstalk, as every MZI switch has phase errors due to imperfect fabrication. Indeed, in a simple $N \times N$ PILOSS topology, only N switches need to be in the cross state, and the remaining $N^2 - N$ switches need to remain in the bar state. Therefore the reduction of the initial phase error is critical for the reduction of the electric power consumption [39].

Fourth, the fabrication imperfections should also be taken into consideration, for example, various kinds of wavelength selective switches (WSSs) have been demonstrated, such as those using ring resonators [86–88], AWGs [89,90], Echelle gratings [91], and contra-directional couplers [19]. However, fluctuations of the waveguide core width and thickness as small as 0.5–1 nm that are unachievable even by an ideal fabrication can result in crosstalk levels up to ~ -20 dB for

WSSs in general [92]. A tandem configuration of these devices is one practical solution to achieve less than −30-dB crosstalk.

Last but not least, because no effective optical memories exist, contention resolution is one of the most critical functionalities that need to be addressed for optical switching nodes. Several approaches have been proposed to resolve contention in OCS:

1. Time domain: the contending packet/burst can be stored in fixed fiber delay line (FDL) or electronic buffer.
2. Wavelength domain: by means of wavelength conversion, the packet/burst can be forwarded to alternative wavelength channels.
3. Space domain: contending burst is forwarded to another output port (deflection routing).

The techniques based on FDL, wavelength conversion and deflection routing all increase significantly the system complexity in terms of routing control and packet synchronization. Moreover, they degrade the power and quality of the signals which would result in limited and fixed buffering time. A promising solution is to exploit the electronic buffer at the edge nodes with optical flow control. Also, to minimize latency, the optical switch should be as close as possible to the edge nodes, and fast decision-making is required. This is feasible in a DC environment with interconnects ranging from few to hundreds of meters.

8.10 Summary

The explosive growth of data traffic in large-scale data centers and HPC systems raises ever-growing performance demands of the interconnection networks regarding bandwidth, cost, power consumption, etc. Optical switching represents a promising solution to address these demands. This chapter focused on the optical switching technologies that can potentially be implemented in the next-generation DCNs. First, the key figures of merit to evaluate optical switches for data center applications were highlighted, for example, the scalability, switching latency, cost, and power efficiency. Afterwards, overviews of typical switch fabric architectures and integrated optical switching technologies were provided, elaborating on their characteristics as well as relevant strengths and shortcomings. The latest optical switch demonstrations based on both silicon and InP photonic integration platforms were presented, with an analysis of their distinct performance properties and suitable application scenarios. Furthermore, hybrid III−V-on-Si optical switches that combine the strong points of both platforms were also discussed. As for practical applications, different data center/HPC optical interconnect schemes powered by optical switches were analyzed with diverse switching performance requirements, such as the interconnection fabric for disaggregated data centers. Finally, the remaining technical obstacles and future development directions for optical switches in data centers were envisioned. In the foreseeable future, it is expected that optical switching technologies will play an increasingly important role in the

data center/HPC interconnect networks, enabling advances in different massive-data applications at lower cost and energy consumption.

References

[1] L. David. IEEE 802.3 Ethernet Working Group Communication. San Diego, 2012. <https://www.readkong.com/page/ieee-802-3-ethernet-working-group-communication-2816783>.

[2] Cisco Visual Networking, The Zettabyte Era—Trends and Analysis. Cisco White Paper, 2017.

[3] T. Barnett, S. Jain, U. Andra, T. Khurana, Cisco Visual Networking Index (VNI) Complete Forecast Update, 2017−2022 Americas/EMEAR Cisco Knowledge Network (CKN) Presentation, p. 125, 2022.

[4] N. Farrington, A. Andreyev, Facebook's data center network architecture, in: Proceedings of the Optical Interconnects Conference, Santa Fe, NM, May 2013, pp. 49−50.

[5] S. Kanev, et al., Profiling a warehouse-scale computer, in: Proceedings of the Forty-Second Annual International Symposium on Computer Architecture, New York, NY, USA, 2015, pp. 158−169.

[6] Q. Cheng, M. Bahadori, M. Glick, S. Rumley, K. Bergman, Recent advances in optical technologies for data centers: a review, Optica 5 (11) (2018) 1354.

[7] K. Sato, Realization and application of large-scale fast optical circuit switch for data center networking, J. Light. Technol. 36 (7) (2018) 1411−1419.

[8] D. Miller, Device requirements for optical interconnects to silicon chips, Proc. IEEE 97 (7) (2009) 1166−1185.

[9] B.G. Lee, Photonic switch fabrics in computer communications systems, in: Proceedings of the Optical Fiber Communication Conference Exposition (OFC), March 2018, pp. 1−22.

[10] S. Rumley, et al., Optical interconnects for extreme scale computing systems|Parallel Computing. New York. <https://dl.acm.org/doi/10.1016/j.parco.2017.02.001>.

[11] K. Wen, et al., Flexfly: enabling a reconfigurable dragonfly through silicon photonics, in: SC'16: Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis, 2016, pp. 166−177.

[12] Q. Cheng, S. Rumley, M. Bahadori, K. Bergman, Photonic switching in high performance datacenters [invited], Opt. Express 26 (12) (2018) 16022.

[13] C.F. Lam, Optical network technologies for datacenter networks (invited paper), in: Proceedings of the Conference on Optical Fiber Communication (OFC/NFOEC), Collocated National Fiber Optic Engineers Conference, March 2010, pp. 1−3.

[14] C. Minkenberg, R. Krishnaswamy, A. Zilkie, D. Nelson, Co-packaged datacenter optics: opportunities and challenges, IET Optoelectron. 15 (2) (2021) 77−91.

[15] X. Zhou, H. Liu, R. Urata, Datacenter optics: requirements, technologies, and trends (invited paper), Chin. Opt. Lett. 15 (5) (2017) 120008.

[16] G. Wang, et al., c-Through: part-time optics in data centers, in: Proceedings of the ACM SIGCOMM 2010 Conference, New York, NY, USA, 2010, pp. 327−338.

[17] Y.-K. Yeo, Q. Huang, L. Zhou, Large port-count optical crossconnects for data centers (invited), in: Proceedings of the International Confeence on Photonics Switching (PS), September 2012, pp. 1−3.

[18] Z. Cao, R. Proietti, S.J.B. Yoo, Hi-LION: hierarchical large-scale interconnection optical network with AWGRs [invited], J. Opt. Commun. Netw. 7 (1) (2015) A97−A105.
[19] J. Perelló, et al., All-optical packet/circuit switching-based data center network for enhanced scalability, latency, and throughput, IEEE Netw. 27 (6) (2013) 14−22.
[20] A. Greenberg, et al., VL2: a scalable and flexible data center network, SIGCOMM Comput. Commun. Rev. 39 (4) (2009) 51−62.
[21] A. Roy, et al., Inside the Social Network's (Datacenter) Network, in: Proceedings of the 2015 ACM Conference on Special Interest Group on Data Communication. San Diego. <https://dl.acm.org/doi/10.1145/2785956.2787472>.
[22] A. Singh, et al., Jupiter rising: a decade of Clos topologies and centralized control in Google's datacenter network, SIGCOMM Comput. Commun. Rev. 45 (4) (2015) 183−197.
[23] G. Kamiya. Data Centres and Data Transmission Networks—Analysis, IEA. Paris. <https://www.iea.org/reports/data-centres-and-data-transmission-networks>.
[24] N. Farrington, et al., Helios: a hybrid. electrical/optical switch. Architecture modular Data Center, 2010, pp. 339−350.
[25] Q. Cheng, A. Wonfor, J.L. Wei, R.V. Penty, I.H. White, Low-energy, high-performance lossless 8×8 SOA switch, in: Proceedings of the Optical Fiber Communication Conference, 2015, p. Th4E. 6.
[26] S. Rumley, Design space exploration of the dragonfly topology, in: High Performance Computing: ISC High Performance 2017 International Workshops, DRBSD, ExaComm, HCPM, HPC-IODC, IWOPH, IXPUG, P^ 3MA, VHPC, Visualization at Scale, WOPSSS, Frankfurt, Germany, June 18−22, 2017, Revised Selected Papers, 2017, vol. 10524, p. 57.
[27] H. Okayama, Y. Okabe, T. Kamijoh, N. Sakamoto, Optical switch array using banyan network, IEICE Trans. Commun. E82-B (2) (1999) 365−372.
[28] H. Wang, A. Wonfor, K.A. Williams, R.V. Penty, I.H. White, Demonstration of a lossless monolithic 16x16 QW SOA switch, in: Proceedings of the Thirty-Fifth European Conference on Optical Communication (ECOC 2009) 20−24 September 2009, Vienna, 2009, pp. 6.2.3−1/2.
[29] H.S. Hinton, An Introduction to Photonic Switching Fabrics, Springer Science & Business Media, 2013.
[30] K. Tanizawa et al., Silicon photonic 32×32 strictly-non-blocking blade switch and its full path characterization, in: Proceedings of the Twenty-First OptoElectronics and Communications Conference (OECC) Held Jointly 2016 International Conference on Photonics Switching (PS), July 2016, pp. 1−3.
[31] M. Yano, F. Yamagishi, T. Tsuda, Optical MEMS for photonic switching-compact and stable optical crossconnect switches for simple, fast, and flexible wavelength applications in recent photonic networks, IEEE J. Sel. Top. Quantum Electron. 11 (2) (2005) 383−394.
[32] J. Kim, et al., 1100 x 1100 port MEMS-based optical crossconnect with 4-dB maximum loss, IEEE Photonics Technol. Lett. 15 (11) (2003) 1537−1539.
[33] A. Wonfor, H. Wang, R.V. Penty, I.H. White, Large port count high-speed optical switch fabric for use within datacenters, J. Opt. Commun. Netw. 3 (8) (2011) A32−A39.
[34] M. Ding, A. Wonfor, Q. Cheng, R.V. Penty, I.H. White, Hybrid MZI-SOA InGaAs/InP photonic integrated switches, IEEE J. Sel. Top. Quantum Electron. 24 (1) (2018) 1−8.
[35] T. Chu, L. Qiao, W. Tang, D. Guo, W. Wu, Fast, high-radix silicon photonic switches, in: Proceedings of the Optical Fiber Communication Conference Exposition (OFC), March 2018, pp. 1−3.

[36] Q. Cheng, et al., Ultralow-crosstalk, strictly non-blocking microring-based optical switch, Photon. Res. 7 (2) (2019) 155−161.
[37] Q. Cheng, M. Bahadori, Y.-H. Hung, Y. Huang, N. Abrams, K. Bergman, Scalable microring-based silicon Clos switch fabric with switch-and-select stages, IEEE J. Sel. Top. Quantum Electron. 25 (5) (2019) 1−11.
[38] T. Shimoe, K. Hajikano, K. Murakami, Path-Independent Insertion Loss Optical Space switch, 1987, p. WB2.
[39] K. Suzuki, et al., Low-insertion-loss and power-efficient 32×32 silicon photonics switch with extremely high-Δ silica PLC connector, J. Light. Technol. 37 (1) (2018) 116−122.
[40] Y. Ma, et al., Ultralow loss single layer submicron silicon waveguide crossing for SOI optical interconnect, Opt. Express 21 (24) (2013) 29374−29382.
[41] K. Suzuki et al., Strictly Non-Blocking Silicon Photonics Switches, 2016, p. AF3G-1.
[42] L. Qiao, W. Tang, T. Chu, 16×16 non-blocking silicon electro-optic switch based on Mach-Zehnder interferometers, in: Proceedings of the Optical Fiber Communication Conference, Paper Th1C.2, March 2016, p. Th1C.2.
[43] Y. Fu, T. Ye, W. Tang, T. Chu, Efficient adiabatic silicon-on-insulator waveguide taper, Photon. Res. 2 (3) (2014) A41.
[44] R. Konoike, H. Matsuura, K. Suzuki, H. Kawashima, K. Ikeda, Polarization-Insensitive Low-Crosstalk 8×8 Silicon Photonics Switch with 9×13.5 cm^2 Control Board, 2020, pp. 1−3.
[45] K. Tanizawa, K. Suzuki, K. Ikeda, S. Namiki, H. Kawashima, Non-duplicate polarization-diversity 8×8 Si-wire PILOSS switch integrated with polarization splitter-rotators, Opt. Express 25 (10) (2017) 10885−10892.
[46] K. Suzuki, et al., Nonduplicate polarization-diversity 32×32 silicon photonics switch based on a SiN/Si double-layer platform, J. Light. Technol. 38 (2) (2020) 226−232.
[47] S. Han, T.J. Seok, K. Yu, N. Quack, R.S. Muller, M.C. Wu, Large-scale polarization-insensitive silicon photonic MEMS switches, J. Light. Technol. 36 (10) (2018) 1824−1830.
[48] S. Nakamura, S. Yanagimachi, H. Takeshita, A. Tajima, T. Hino, K. Fukuchi, Optical switches based on silicon photonics for ROADM application, IEEE J. Sel. Top. Quantum Electron. 22 (6) (2016) 185−193.
[49] L. Qiao, Researches on the Large-Scale Silicon Optical Switch Array. Ph.D. thesis, University of Chinese Academy of Sciences, 2017.
[50] L. Qiao, W. Tang, T. Chu, 32×32 silicon electro-optic switch with built-in monitors and balanced-status units, Sci. Rep. 7 (1) (2017) 1−7.
[51] I. Armstrong, I. Andonovic, A. Kelly, Semiconductor optical amplifiers: performance and applications in optical packet switching [Invited], J. Opt. Netw. 3 (12) (2004) 882−897.
[52] Q. Cheng, A. Wonfor, R.V. Penty, I.H. White, Scalable, low-energy hybrid photonic space switch, J. Light. Technol. 31 (18) (2013) 3077−3084.
[53] R. Stabile, A. Albores-Mejia, K. Williams, Monolithic active-passive 16×16 optoelectronic switch, Opt. Lett. 37 (22) (2012) 4666−4668.
[54] M. Smit, et al., <https://iopscience.iop.org/article/10.1088/0268-1242/29/8/083001/meta>.
[55] S.C. Nicholes, M.L. Mašanović, B. Jevremović, E. Lively, L.A. Coldren, D.J. Blumenthal, The world's first inp 8x8 monolithic tunable optical router (MOTOR) operating at 40 Gbps line rate per port, in: Proceedings of the Optical Fiber Communication Conference; National Fiber Optic Engineers Conference, San Diego, California, 2009, p. PDPB1.
[56] X. Ye, Y. Yin, S.B. Yoo, P. Mejia, R. Proietti, V. Akella, DOS-A scalable optical switch for datacenters, 2010, pp. 1−12.

[57] J. Gripp, et al., Optical switch fabrics for ultra-high-capacity IP routers, J. Light. Technol. 21 (11) (2003) 2839.
[58] R. Stabile, A. Rohit, K.A. Williams, Monolithically integrated 8 × 8 space and wavelength selective cross-connect, J. Light. Technol. 32 (2) (2014) 201–207.
[59] R. Stabile, A. Rohit, K.A. Williams, Dynamic multi-path WDM routing in a monolithically integrated 8 × 8 cross-connect, Opt. Express 22 (1) (2014) 435–442.
[60] Q. Cheng, et al., First demonstration of automated control and assessment of a dynamically reconfigured monolithic 8 × 8 wavelength-and-space switch, J. Opt. Commun. Netw. 7 (3) (2015) A388–A395.
[61] R. Stabile, K. Williams, Optical Power Meter Co-Integrated With a Fast Optical Switch for on-Chip OSNR Monitoring, 2015, pp. 360–362.
[62] T. Matsumoto, et al., In-line optical amplification for Si waveguides on 1 × 8 splitter and selector by flip-chip bonded InP-SOAs, in: Proceedings of the Optical Fiber Communication Conference Exhibition (OFC), March 2016, pp. 1–3.
[63] T. Matsumoto, et al., Hybrid-integration of SOA on silicon photonics platform based on flip-chip bonding, J. Light. Technol. 37 (2) (2019) 307–313.
[64] R. Konoike, et al., SOA-integrated silicon photonics switch and its lossless multistage transmission of high-capacity WDM signals, J. Light. Technol. 37 (1) (2018) 123–130.
[65] R. Konoike, et al., Gain-integrated 8 × 8 silicon photonics multicast switch with on-chip 2 × 4-ch. SOAs, J. Light. Technol. 38 (11) (2020) 2930–2937.
[66] T. Aalto, et al., Open-access 3-μm SOI waveguide platform for dense photonic integrated circuits, IEEE J. Sel. Top. Quantum Electron. 25 (5) (2019) 1–9.
[67] N. Tessema, et al., Wavelength selective photonic integrated switches for ROADM node functionality in ultrahigh capacity metro network, in: Proceedings of the International Conference on Optical Network Design and Modelling (ONDM), Gothenburg, Sweden, June 2021, pp. 1–5.
[68] D. Tanaka, et al., Ultra-small, self-holding, optical gate switch using Ge2Sb2Te5 with a multi-mode Si waveguide, Opt. Express 20 (9) (2012) 10283–10294.
[69] K. Kato, M. Kuwahara, H. Kawashima, T. Tsuruoka, H. Tsuda, Current-driven phase-change optical gate switch using indium–tin-oxide heater, Appl. Phys. Express 10 (7) (2017) 072201.
[70] K. Shibuya, Y. Atsumi, T. Yoshida, Y. Sakakibara, M. Mori, A. Sawa, Silicon waveguide optical modulator driven by metal–insulator transition of vanadium dioxide cladding layer, Opt. Express 27 (4) (2019) 4147–4156.
[71] T. Murai, Y. Shoji, N. Nishiyama, T. Mizumoto, Nonvolatile magneto-optical switches integrated with a magnet stripe array, Opt. Express 28 (21) (2020) 31675–31685.
[72] M. Naumov, et al., Deep Learning Training in Facebook Data Centers: Design of Scale-Up and Scale-Out Systems, arXiv:2003.09518 [cs], August 2020. <http://arxiv.org/abs/2003.09518>.
[73] Z. Zhu, et al., Distributed deep learning training using silicon photonic switched architectures, APL Photonics 7 (3) (2022) 030901.
[74] Z. Zhu, et al., Photonic switched optically connected memory: an approach to address memory challenges in deep learning, J. Light. Technol. 38 (10) (2020) 2815–2825.
[75] C. Minkenberg, et al., Performance Benefits of Optical Circuit Switches for Large-Scale Dragonfly Networks, 2016, pp. W3J–W33.
[76] R. Yasudo, J. Coutinho, A. Varbanescu, W. Luk, H. Amano, T. Becker, Performance Estimation for Exascale Reconfigurable Dataflow Platforms, 2018, pp. 314–317.

[77] G. Massari, A. Pupykina, G. Agosta, W. Fornaciari, Predictive Resource Management for Next-Generation High-Performance Computing Heterogeneous Platforms, 2019, pp. 470−483.
[78] Y. Shen, M.S. Glick, K. Bergman, Silicon photonic-enabled bandwidth steering for resource-efficient high performance computing, 2019, 10946, pp. 17−25.
[79] F. Yan, W. Miao, O. Raz, N. Calabretta, Opsquare: a flat DCN architecture based on flow-controlled optical packet switches, J. Opt. Commun. Netw. 9 (4) (2017) 291−303.
[80] X. Guo, X. Xue, F. Yan, B. Pan, G. Exarchakos, N. Calabretta, DACON: a reconfigurable application-centric optical network for disaggregated data center infrastructures [invited], J. Opt. Commun. Netw. 14 (1) (2022) A69−A80.
[81] S. Namiki, R. Matsumoto, T. Inoue, K. Ishii, Scaling Silicon Photonic Switches and Digitalization of the Optical Layer, 2021, pp. W1C−W11.
[82] R. Matsumoto, et al., Fully-Loaded Operation of 0.29-pJ/bit Wall-Plug Efficiency, 81.9-Tb/s Throughput 32 × 32 Silicon Photonics Switch, 2021, pp. Tu6A-2.
[83] M. Takahashi, Y. Uchida, S. Yamasaki, J. Hasegawa, T. Yagi, Compact and low-loss coherent mixer based on high Δ ZrO 2-SiO 2 PLC, J. Light. Technol. 32 (17) (2014) 3081−3088.
[84] J. Hasegawa, et al., 32-port 5.5%-Δ Silica-Based Connecting Device for Low-Loss Coupling Between SMFs and Silicon Waveguides, 2018, pp. 1−3.
[85] L. Chen, Y. Chen, Compact, low-loss and low-power 8 × 8 broadband silicon optical switch, Opt. Express 20 (17) (2012) 18977−18985.
[86] P. Prabhathan, Z. Jing, V.M. Murukeshan, Z. Huijuan, C. Shiyi, Discrete and fine wavelength tunable thermo-optic WSS for low power consumption C + L band tunability, IEEE Photonics Technol. Lett. 24 (2) (2011) 152−154.
[87] N. Eid, R. Boeck, H. Jayatilleka, L. Chrostowski, W. Shi, N.A. Jaeger, FSR-free silicon-on-insulator microring resonator based filter with bent contra-directional couplers, Opt. Express 24 (25) (2016) 29009−29021.
[88] A.S. Khope, et al., Multi-wavelength selective crossbar switch, Opt. Express 27 (4) (2019) 5203−5216.
[89] C.R. Doerr, L.L. Buhl, L. Chen, N. Dupuis, Monolithic flexible-grid 1 × 2 wavelength-selective switch in silicon photonics, J. Light. Technol. 30 (4) (2012) 473−478.
[90] F. Nakamura, et al., Integrated silicon photonic wavelength-selective switch using wavefront control waveguides, Opt. Express 26 (10) (2018) 13573−13589.
[91] T.J. Seok, et al., Silicon photonic wavelength cross-connect with integrated MEMS switching, APL Photonics 4 (10) (2019) 100803.
[92] K. Okamoto, Wavelength-division-multiplexing devices in thin SOI: advances and prospects, IEEE J. Sel. Top. Quantum Electron. 20 (4) (2013) 248−257.

This page intentionally left blank

Photonic-integrated circuits and optical fabrics for heterogeneous computing systems

9

Doug Carmean[1], Robert Stone[1], Mark Wade[2], S.J. Ben Yoo[3], Ram Huggahalli[4] and Larry Dennison[5]

[1]Meta Platforms Inc., Menlo Park, CA, United States, [2]Ayar Labs, Santa Clara, CA, United States, [3]Department of Electrical and Computer Engineering, University of California Davis, Davis, CA, United States, [4]Microsoft Corporation, One Microsoft Way, Redmond, WA, United States, [5]NVIDIA Corporation, Santa Clara, CA, United States

9.1 Problem statement

Modern computing workloads including artificial intelligence, machine learning, and graph analytics are facing severe challenges related to data movement in latency, energy consumption, and productivity. The speed gap between processing units and memory—the "memory wall" [1] problem—has become a critical system performance bottleneck. It is also well known that application-specific computing systems optimally designed and configured for a given workload offer much higher (often 50× or more) [2,3] energy efficiency and throughput compared to general-purpose computing systems. Heterogeneous computing systems (HCS) exploiting energy and performance benefits of combining different domain-specific processor architectures have emerged in modern computing systems of different scales—ranging from Systems-on-chip to data center and HPC systems. Increasing numbers of processing and memory resources have led to HCSs of increasing scale, with significant computing power potential. On the other hand, application domains like graph analytics and machine learning now process terabyte data sets which are growing to petabytes. Therefore the performance gains of running these workloads in HCSs highly depend on the scaling of the memory architecture. In an HCS, some accelerators may have embedded memory dies and some others may rely on the memory connected to the system bus. This nonuniform memory architecture significantly increases the complexity of programming such heterogeneous systems. Ideally, we want any compute node to be able to communicate to any memory node with sufficient bandwidth without introducing complex programming models. We want to allow programmers to reference any data from any device. Thus it is desirable that the HCS should reconfigure its processor (i.e., CPUs, GPUs, ASICs) to memory (volatile and nonvolatile memories) interconnection, following the change in the dataflow to optimize its throughput and energy efficiency across their disaggregated system.

Unfortunately, current systems utilizing electrical interconnects face several challenges for implementing such HCSs. First, the fundamental limits imposed by the laws of physics, such as impedance, skin effect, electromagnetic interference,

Integrated Photonics for Data Communication Applications. DOI: https://doi.org/10.1016/B978-0-323-91224-2.00004-7
© 2023 Elsevier Inc. All rights reserved.

and dispersion, are limiting the bandwidth and the unrepeated distance of electronic interconnects. This problem gets worse as the electrical wire cross-section gets smaller. In addition, the number of connection pins is physically limited, forcing higher bandwidth per pin. Struggling with this *Pin Wall*, electrical links do not provide enough bandwidth-distance product for HCSs, and have high latency when accessing remote memory which incurs performance, energy, and programmability overheads. Secondly, the accelerating trends for higher-throughput, lower-latency, lower-power, heterogeneity, and disaggregation requirements in HCSs are no longer compatible with traditional architectures that rely on large electronic controllers and electronic crossbar switches interconnecting memory and processors. Thirdly, the lack of flexibility of electrical interconnects as well as their distance-dependent energy consumption make an energy-efficient design of HCS challenging, unable to disaggregate the computing resources in a heat density limited electrical platform.

Photonic interconnect fabrics offers a breakthrough in disaggregating the computing resources and overcoming the data movement challenges addressing energy efficiency, throughput, reconfigurability, productivity, and heterogeneity of computing for various workloads. This Chapter covers innovative solutions enabled by distance-independent interconnects at low power at essentially unlimited bandwidth compared to electronic counterparts. When combined with wavelength-division multiplexing (WDM) and silicon photonics, optical interconnects become disruptive and promising technologies for future HCSs. Naeemi [4] have defined a partition length (above which information is more efficiently transported by photons rather than electrons), and Ref. [5] summarizes this in Fig. 9.1. Note that above 20 GHz, the partition length

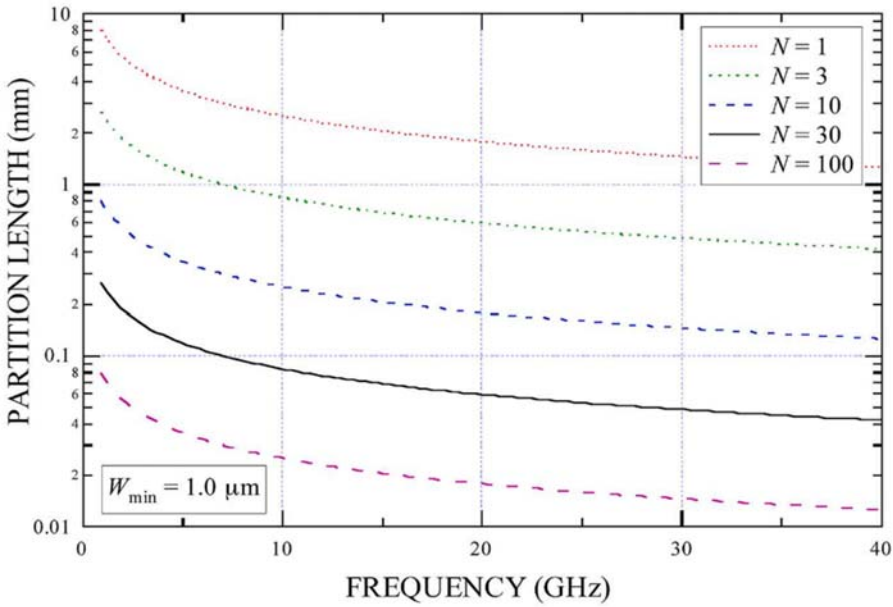

Figure 9.1 The partition length given for a wire/waveguide width of 1 μm. Note that above 20 GHz, the partition length for a system using 30 wavelength channels is less than 60 μm.

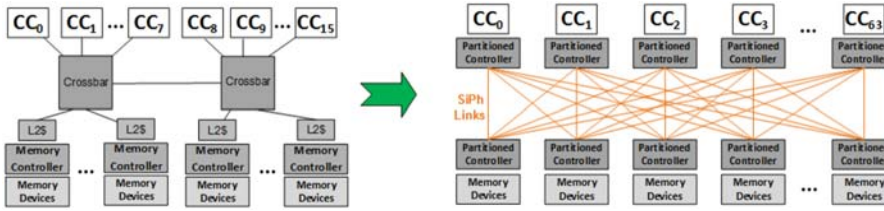

Figure 9.2 (Left) Overview of baseline memory system where different core clusters (*CCs*) share a crossbar, a single read/write queue per channel, and a last level cache. (Right) Proposed memory system addresses the contention by providing dedicated queues for each core cluster to send memory request to every channel through an all-to-all interconnect.

for a system using 30 wavelength channels is less than 60 micro-meters. Secondly, WDM photonic interconnects offer scalable, high-throughput, and low-latency interconnects between a very large number (e.g., over 524,000 for 3D Hyper-X LION [6]) of distributed/disaggregated compute nodes. Fig. 9.2 shows the contrast between the current method (Fig. 9.2, left) and the photonics-enabled method (Fig. 9.2, right). The large and centralized electronic crossbar switches with electronic *I/Os* with centralized controllers are replaced by small and distributed electronic switches and partitioned controllers with parallel and contentionless parallel photonic interconnect (orange). When utilizing WDM and silicon photonic interconnect technologies, orders-of-magnitude improvements in latency, throughput, and energy efficiency can be achieved for data movements in the HCSs. Third, optical reconfigurability and bandwidth steering in the HCSs can help achieve the best match between the interconnection topology and the expected data flow for any given application. Fig. 9.3 shows communication patterns in three application examples. While the hardwired electronic interconnects cannot be pruned or reconfigured easily, the optical wavelength assignments allow bandwidth steering by 100-fold in the 100 wavelength systems, enabling support of diverse and sparse applications with programmable hardware platforms. Lastly, silicon photonic-integrated circuits (PICs) offer an ideal fabric platform for electronic–photonic codesign, cointegration, and copackaging within the same ecosystem. The traditional electronic wafer manufacturing and packaging foundries are already commercializing electronic–photonic design automation, multiproject-wafer runs, and test-and-packaging capabilities.

9.2 Applications space with requirements

As noted above, a key consideration in scaling device and system performance is the evolution of available CMOS technologies. This is summarized in Fig. 9.4, based on the models and roadmap data from The International Roadmap for Device and Systems [7]. Here we follow the conventional naming convention for process nodes, the time axis of this data represents approximate commercial availability of a given technology. Three key device performance parameters are included and are

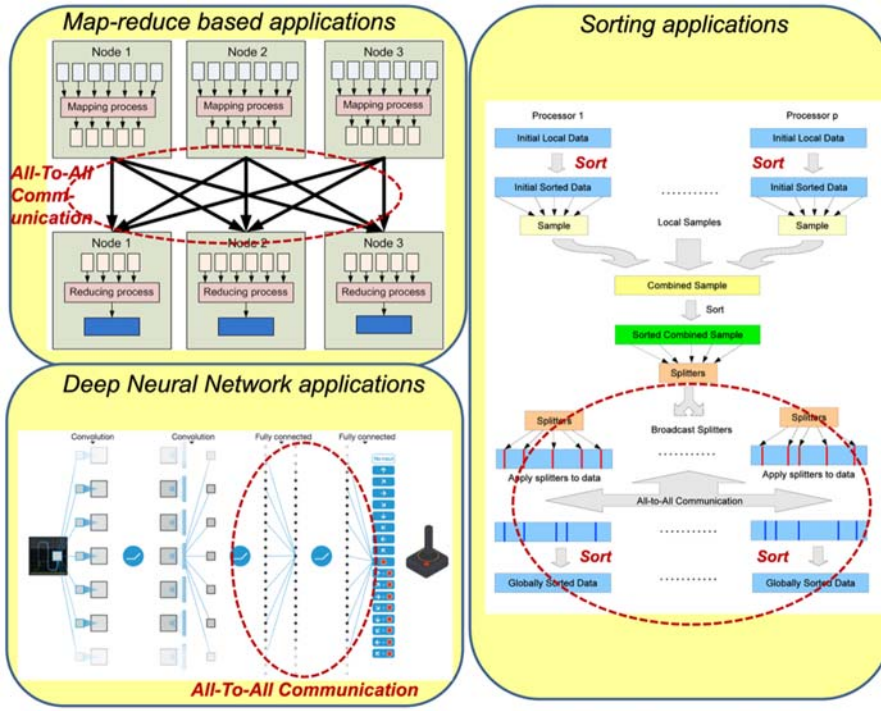

Figure 9.3 Examples of applications showing very different data movement patterns. Map-reduce, deep neural network, and sorting applications all show all-to-all communication pattern at some point of the workflow but not throughout the entire workflow.

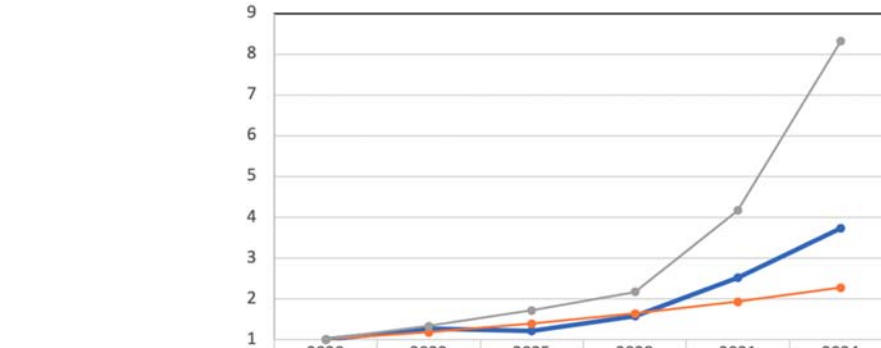

	2020	2022	2025	2028	2031	2034
CMOS Technology Node Class	5	3	2.1	1.5	1	0.7
Relative Power Density	1.0	1.3	1.2	1.6	2.5	3.7
Relative IO Density	1.0	1.2	1.4	1.6	1.9	2.3
Relative Logic Density	1.0	1.3	1.7	2.2	4.2	8.3

Figure 9.4 Forecast CMOS technology performance scaling.

normalized to the 5 nm process node; power density, IO density, and logic density. As can be observed in the data, although logic density is forecast to scale consistently with introduction of subsequent node, IO density is not anticipated to scale at the same rate. There are likely several factors involved in this scaling discrepancy between logic and IO related to not only limitations in packaging technology (i.e., minimum practical device bump pitch, currently in the order of 40 μm), but also the fact that in general the analog circuit area for an IO block does not follow the same area scaling as logic cells. Further magnifying this difference is the fact that subsequent generations of device place higher bandwidth demand on the IO (whether lower speed IO to memory, or serialized higher speed interconnects to other system elements). Although the density of lower speed IO tends to be limited by the pitch of the available IO bumps, high speed IO density typically is circuit area limited, with higher lane rate SERDES requiring more complex equalization and associated link layer logic.

Although the scaling of both IO and logic density enables advancements in performance, the increasing power density also presents a problem for the system designer. Current leading-edge compute and NPU devices are already exceeding 500 W per device. Thermal and packaging solutions are today primarily reliant upon a radiative air cooled heatsink with either heat pipes or vapor chambers to improve heat transfer from the device heat spreader. Ambient forced air is pulled into the system via front panel perforations, and ducted over the heatsink using an array of fans in the rear of the system. To support higher power densities, techniques, such as those already demonstrated in HPC systems, using liquid cooling will likely be required, following the forecast trend of a 4× power density increase, as a similar sized scaled die would dissipate approximately 2 kW. As well as the thermal challenges associated with such future large area devices, the anticipated lower core voltage ($\sim$0.5 V) of these process nodes makes the power supply design and distribution more difficult as several thousands amps of current are required.

Two major technology barriers are thus identified to continuing to scale current device architectures, even given advances in process technology; cooling, and power delivery. These barriers to continued device scaling suggest spatial separation (or disaggregation) of large single devices into a greater number of smaller subsystems, lowering the overall power density and thus relaxing the demands on power delivery and cooling.

System disaggregation however introduces a new set of challenges. As can be seen in Fig. 9.5 which summarizes interconnect efficiency for various reach technologies, efficiency degrades as reach is increased. For on-chip and in-package IO (such as internal chip busses and HBM memory interfaces), simple low power analog circuits are used to drive low loss electrical channels at a relatively low lane rate. Intrachip CMOS wiring or advanced packaging (such as Si interposers, Intel's EMIB or TSMC's InFO) provide fine pitch wiring geometries which enable practical instantiation of high bandwidth links by aggregation of multiple lanes into a high bandwidth bus at low relative cost. However, as the interconnect distance increases to on-board, or within the rack distances, designers are forced to aggregate to a fewer number of lanes, as the PCB and electrical cabling wiring densities cannot support the same

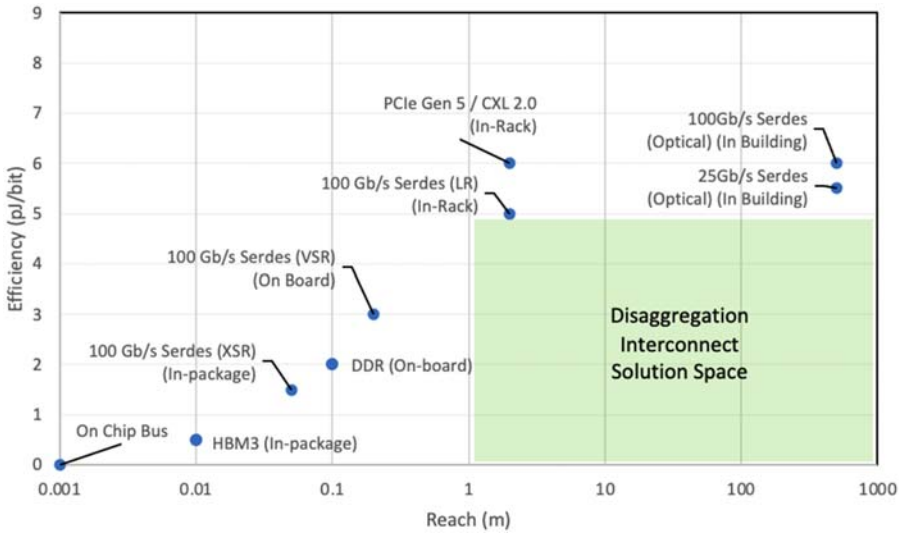

Figure 9.5 Interconnect efficiency versus reach.

number of lanes. This requiring serialization to higher lane rates. Higher speed, longer reach SerDes (such as 100 Gb/s in Fig. 9.5) require greater equalization capability to compensate for the higher frequency material losses, reflections and other channel nonidealities. This in-turn lowers the overall efficiency of the interconnect. As the lane rate increases the efficiency of the interconnect is reduced, however the reach is increased. System disaggregation therefore leads to a trade-off between bandwidth and reach within a manageable power envelope per device.

The enabling technology to continue to scale system performance in the postdisaggregation regime is self-evident; a high bandwidth, high density and efficient optical interconnect. Current optical solutions are encumbered by the dominance of the electrical SerDes, and due to the identified lower efficiency (and cost) are used by system designers as a last resort alternative when the electrical interfaces can no longer provide adequate reach. To address this efficiency gap, codesign of the host ASIC and optical PICs will be required, with 3D integration to eliminate electrical loss between the ASIC and optics (see chapters 1 and 7 for further information).

9.3 Interfaces

Efficient electrical and optical interfaces are critical building blocks for optical fabrics. Fig. 9.6 illustrates the anticipated evolution of both electrical and optical interfaces with the assumption of appropriate transceiver and connector technology in between. In any system, the choice of copper or optical connectivity has primarily been a distance versus data rate trade-off. At a specific combination of aggregate link bandwidth, data rate, distance targeted by a system, copper can be perceived to exceed a specific system's cost, power, and cable management targets. The thresholds can be

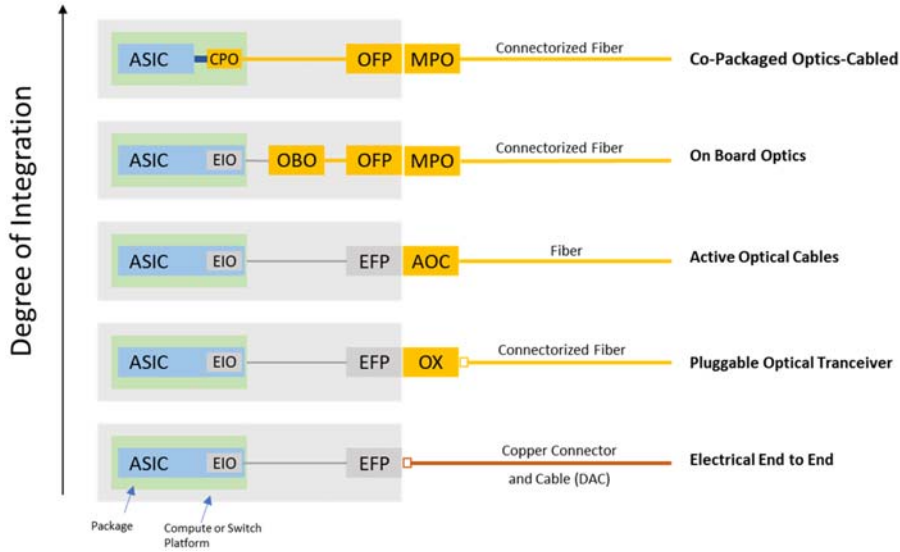

Figure 9.6 Evolution of optical integration.
Source: Optical I/O Chiplets Eliminate Bottlenecks and Unleash Innovation (ayarlabs.com).

implementation specific but have been sufficient to justify optical connectivity pervasively in large-scale systems for interrack connectivity, and the overall trend is to move the electrical-to-optical conversion closer to the ASIC as performance increases.

The separation of the optical engine from the face-plate connector enables two fundamental improvements: (1) the reduction of electrical distance between the SoC (or ASIC) and the optical engine and (2) the simplification of face-plate connector into a passive component. on-board optics (OBO) and copackaged optics (CPO) are steps towards minimizing the electrical distance between the ASIC and the optical engine. When the ASIC and the optical engine are on the same package, the short electrical distance must be supported by electrical layer technology and substrate technology to achieve the best bandwidth density and energy efficiency.

The remainder of the section explores the trade-offs on both the electrical and optical interfaces and envisions the best practical implementations with 2.5 and 3D packaging technologies. We first identify packaging options and their trade-offs.

Organic substrates offer low-cost, medium-density connectivity ($\sim$130 um pitch) for target applications where XSR distance (up to 50 mm) is required for chip layout. Silicon interposers offer higher density connectivity ($<$50 um pitch) but results in a more expensive package and are limited in size. Silicon bridging [cite EMIB, TSMC] embeds high-density connectivity silicon into organic substrates to optimize the performance, cost, and scalability of the package. Organic interposers and hybrid solutions that include silicon bridges are being explored in R&D to further optimize performance and cost. Finally, 3D stacking technologies are beginning to be used in high-end CPU/GPU products. Fig. 9.7 illustrates these options.

The terminology used by industry for interface distances is provided in Table 9.1 as a reference. In forward-looking heterogenous system designs and package

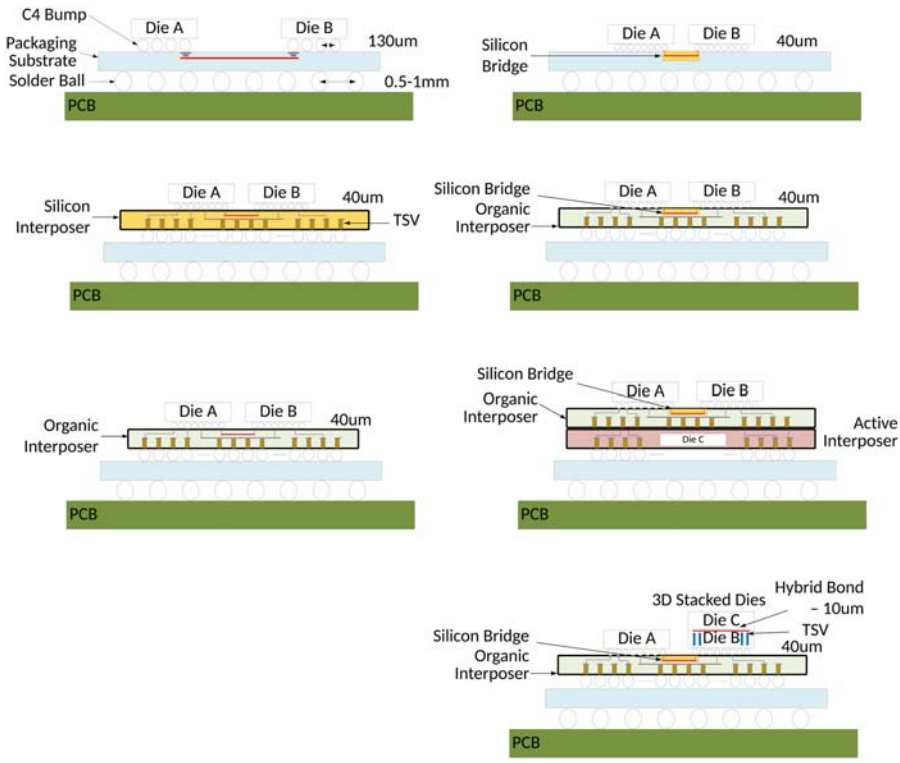

Figure 9.7 Package substrate technologies.

Table 9.1 Electrical interface reach terminology and key properties.

Interface	Application	Length	Loss
USR Ultrashort reach	Bump-to-bump Inside MCP or 3D Stack	<10 mm (0.4 in)	1.5 dB @ 14 GHz 3 dB @ 28 GHz
XSR/LPPI Extremely short reach	Ball-to-Ball Across PCB	<50 mm (2.0 in)	4 dB @ 14 GHz 10 dB @ 28 GHz
VSR/C2M Very short reach	Ball-to-Ball	<200 mm (7.9 in)	10 dB @ 14 GHz 20 dB @ 28 GHz

Source: Optical I/O Chiplets Eliminate Bottlenecks and Unleash Innovation (ayarlabs.com).

implementations, we are focused on USR (<10 mm) and XSR (<50 mm) distances when exploring tight integration of optical engines with ASICs, although we note that some configurations will require VSR (<200 mm) distances.

A summary of recent industry proposals is provided in Table 9.2. The first column: PCIe Gen5, board-level, VSR interface is a reference for comparison with

Table 9.2 Electrical layer research targets.

Substrate	**PCB**	**Organic substrate**				**Si interposer/Si bridge**				**Research targets**	
Distance	**VSR**	**XSR**				**USR**				**<10 mm**	**<50 mm**
Interconnect	**PCIe Gen5**	**PCIe Gen6**	**XSR NRZ**	**XSR PAM4**	**BoW (Fast)**	**BoW (Micro)**	**AIB Gen2**	**UCIe**	**OpenHBI**		
Standard	PCI SIG	PCI SIG	OIF	OIF	ODSA	ODSA	ChipsAlliance	UCIe	ODSA/ JEDEC		
Serial/parallel	Serial	Serial	Serial	Serial	Ser/par	Parallel	Parallel	Parallel	Parallel		
Bump pitch (μm)	500−1000		110−130	110−130	130	40	55	25−55	40		
Data rate (Gbps)	32	64	56	112	16	5	6.4	4−32	10.0	32	80
Reach (mm)	75−125	75−125	30−50	30−50	50	5	5	2−25	4	2	20
Latency	10 ns	<20 ns	10 ns	18−22.5 ns (w/FEC)	5 ns	5 ns	0.9 ns	2 ns	4 ns		
BER (pre-FEC)	10^{-12}	10^{-12}	10^{-12}	50 mm:10^{-9} XXX 20 mm: 10^{-15}	10^{-15}	10^{-15}	10^{-15}	10^{-10}	10^{-15}		
Energy efficiency (*B*), pJ/bit	6	6	1−1.5	0.9−1.4	0.55	0.5	0.5	0.5	0.5	0.4	1.25
Linear BW density (*A*) Gbps/mm	170	170	600	1000	850	1000	1600	1317	1750	6400	1900
Areal BW density Gbps/mm^2	170	170	480	500	680	1000	∼1600	1350	2100		
FOM (*A*/*B*)	28	28	400−600	700−1100	1540	2000	3200	TBD	3500	16000	1520
Cost	Lowest	Lowest	Low	Low	Low	High	High	TBD	High		

current and upcoming package-level electrical layer solutions. Existing XSR interfaces are approximately 5× better in energy efficiency and ~4−6× better in linear bandwidth density. Emerging options with organic substrates include the ODSA proposed BoW-Fast and BoW-Basic (not shown in table) options. BoW-Basic primarily focuses on improving energy efficiency while BoW-Fast uses termination to improve signal integrity and drive a higher data rate for 5× better linear bandwidth density and 10× better energy efficiency relative to a PCB baseline.

At higher cost and by limiting reach, silicon interposer and silicon bridge options can deliver very high bandwidth density. UCIe, BoW-Micro, AIB Gen2 and OpenHBI are emerging industry proposals exceeding 1 Tbps/mm. Optical transceivers that use these physical protocols could operate at the native frequency, or multiplex multiple channels of data onto a higher frequency optical connection. The much better research target indicates that further lowering the distance may permit much higher data rates, higher bandwidth density and FOM. Data rates between 16−32 Gbps/wire will be actively explored in the near future. The last two columns provide guidance on research targets that are deemed to be aggressive but achievable for future interfaces to serve as a comparison point.

The two categories of options, 10−50 mm, organic substrate, and <5 mm, interposer/bridge, should be viewed as complementary and can together enable a wide range of systems applications.

Optical interface design considerations are described in Section 9.8. A critical aspect of scaling bandwidth at low cost is the ability to scale wavelengths. The 100-G lambda multisource agreement (MSA) has specified FR4 as cost-effective way of scaling 100-Gb/s data rates with 4 wavelengths [8]. The FR4 optical lane format has been adopted for 400-Gb/s products with industry beginning to investigate greater wavelengths. The CWWDM MSA (Continuous-Wave Wavelength Division Multiplexing MSA) was formed to propose standards for 8, 16, and 32 wavelengths in O-band [9] to support emerging microring-based technologies and products [10].

Any number or type of ASICs may be connected by a fabric based on optical modules that use a combination of electrical and optical interface options described in the previous section. Recent developments in applying CPO to Ethernet switches (see full discussion in Chapter 7) is motivated by power reduction in Ethernet switch platforms while also improving airflow and reducing the complexity of boards. In the context of future heterogenous systems, we examine the basic connectivity requirements of compute end-points and the potential benefits of using emerging electrical and optical interfaces. The compute platform may be any combination of general-purpose compute and specialized compute elements.

Fig. 9.8 compares an Electrical Baseline where two end-points A and B are connected through a switch completely with standard PCIe serial interfaces. We assume a simple application where the end-points only have a ×16 PCIe Gen5 (64 GB/s) requirement to access a shared resource. Aligned with industry directions, this physical interface may be used by a compute end-point to access a memory end-point through the CXL protocol running over PCIe physicals and through a CXL switch. The memory end-point may be assumed to be part of a pool of

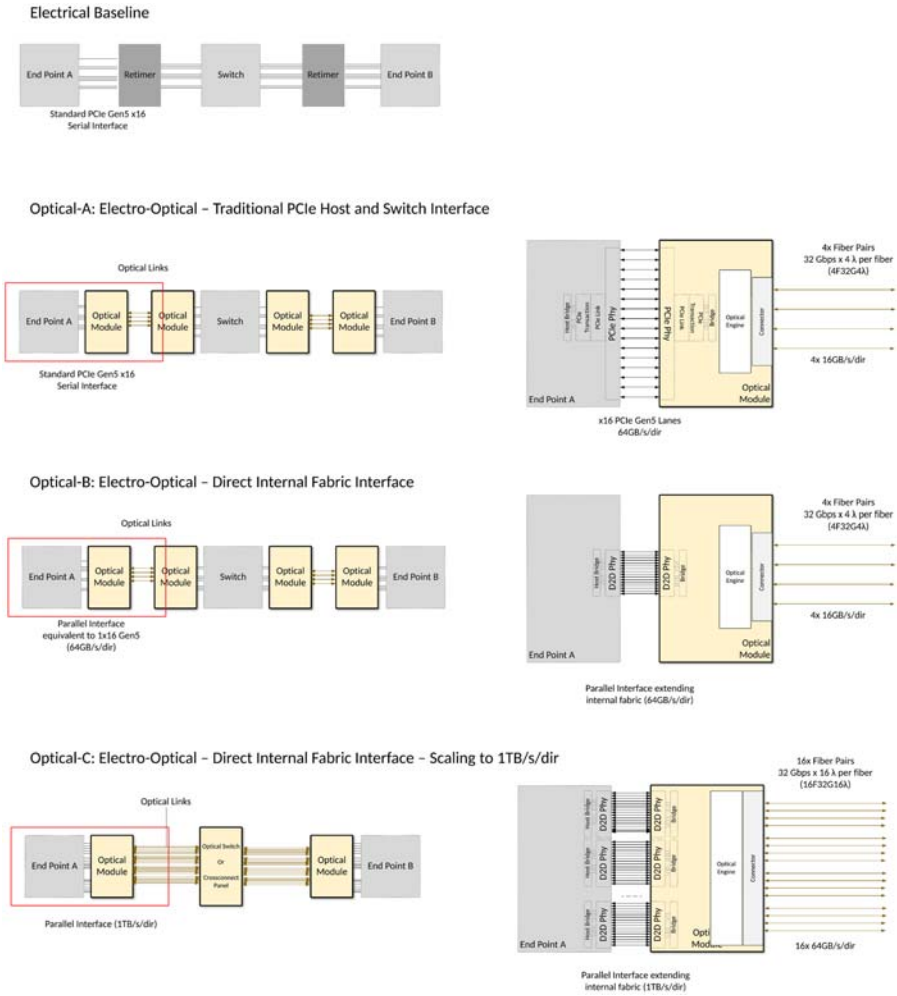

Figure 9.8 Optical module approaches.

memory resources and is disaggregated with the help of retimers to ensure signal strength at 0.5−1 m distance. In this scenario, in addition to the latency of the electrical switch, the primary latency contributors are the PCIe link and physical layer latencies that are added at the entry and exit of every electrical component.

The optical options shown include:

Optical A: the compute node adheres to a standard PCIe SerDes and uses an optical module to convert from 32-Gbps NRZ data rate on copper to optical. The optical module uses 4 fiber pairs with 4 wavelengths to emulate 16 serial lanes.

Optical B: the compute node uses a <5-mm reach parallel electrical interface to the optical module. This has the positive effect of eliminating typical PCIe

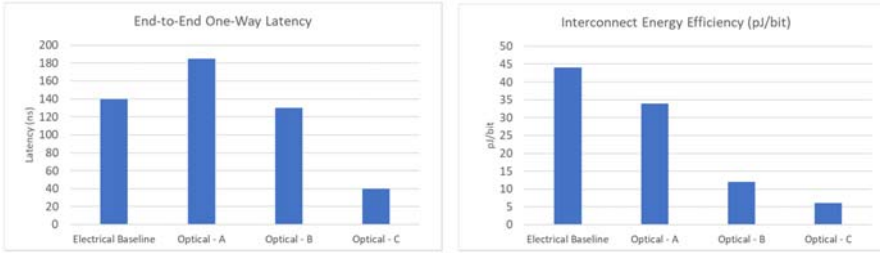

Figure 9.9 Latency and energy efficiency comparison of electrical and optical options.

transaction, link, and physical layers power and latency with potential for a marked improvement as shown in Fig. 9.9.

Optical C: the compute node fully utilizes the bandwidth density of <5 mm reach parallel electrical interface technologies to scale to 1TB/s capability (equivalent to 16, $\times 16$ PCIe Gen5 bandwidth). Such bandwidth is prohibitively expensive with a traditional, VSR approach. On the optical side, the optical module supports 16 fiber pairs with 16 wavelengths with 32 Gbps per wavelength for an aggregate bandwidth of 1 TB/s/dir. In addition to delivering such bandwidth, Fig. 9.9 illustrates the potential for the lowest latency and energy per bit by taking advantage of the 16-way fan-out and the elimination of electrical switch.

At TB/s bandwidth levels, we expect the cost of optics to be much more affordable than electrical options (see Section 9.6).

The transition to optical *I/O* requires a joint optimization between electrical *I/O*, packaging technologies, optical/photonics technologies, and overall communications architectures. In this section, we presented the basics of each of these elements while focusing on key messages, but there are additional layers of complexity to be considered for specific use cases. We expect there will be a multitude of solutions that exist in industry that are optimized for various applications, but the driving trends of advanced packaging, ultraefficient electrical interfaces, and dense energy-efficient optical interfaces will be the focus.

9.4 System level reliability, resiliency, BER, lifetime

A full discussion of component level reliability is presented in Chapter 13. Today's 100-Gbps pluggable modules have a FIT rate of ~ 300. Using a purely linear extrapolation, a 3.2-Tbps module with the same technology would likely have a FIT rate of ~ 9600. System integrators and Cloud Service Providers have identified a target that 3.2-Tbps modules have an FIT rate below 60 to support system annualized failure rates of $<1\%$ in 2023. This section explores whether the goal of 60 FIT for a 3.2-Tbps module is realistic.

In the 1990s CPUs and DRAMs were increasing shipping volumes rapidly. At that time, FIT rates for commodity CPUs were typically far in excess of 1000 FIT.

Failure models were highly conservative in that they were based on accelerated temperature measurements and basic assumptions that every transistor failure resulted in a device failure. Over time, the models were improved with empirical data and the relationship between localized transistor/bit failures and device failures became more nuanced. Importantly, this work required market demand that drove high-volume manufacturing (HVM).

HVM both facilitated correlating low level technology failures with higher level system and component failures and drove the requirement of lower FIT rates. In many cases, the improvements came through sequentially solving problems (onion peeling) that were masked by other problems.

Integrated device manufacturers drove FIT rates by developing test chips and structures that were used to drive both yield improvements and FIT characterization. Intel, for example, developed custom SRAM chips that were used to bring up new process technology and to characterize FIT rates with regular structures that were easy to test and localize failures that occurred during burn-in and life tests. These test chips were critical for both yield improvement and for improving the FIT rates of CPUs that were eventually built on those technology nodes.

Today's 100-Gbps pluggable modules are discrete systems with multiple components assembled on a printed circuit board, housed in small chassis with both electrical and optical connectors. The failure rates of today's modules are dominated by connector and laser failures. While the move to copackaging optics (CPO) will improve the module reliability, the connector and laser reliability will continue to be separate concerns.

Multiple optical vendors have declared support for a goal of 60 FIT for a 3.2 Tbps module. While the approaches to reliability differ and many of the specific techniques to reduce FIT rates remain under development, the goal appears sound. Both the target FIT rate and a systematic capability to measure the failure rates at the component level will be important to inform the system level FIT rates and to drive continued improvement in the FIT rates over time.

It is important to understand the impact of a failure, in addition to having a quantified FIT rate target. For example, a 3.2 Tbps module will likely be constructed of 32×112 Gbps parallel channels. A failure in a single connector, or in a single transceiver will, likely, not impact the other channels on that module. Importantly, the module capabilities are degraded, but the overall functionality is maintained. There are many failure modes that would impact the entire module and likely have a significant impact at the device and system level. Eventually, the failure rates should include some way of describing the impact, perhaps in terms of a blast radius or other factor that would allow the system level FIT calculation to determine if there is a functional or degraded mode impact.

The system level failure rate is a combination of all of the components that are used to build the system. As an example, the data center annualized failure rate for optical interconnect is a function of the modules, connectors, lasers, and switches that are used to form an optical link. Importantly, the failure rate is a function of the topology and capacity of the interconnect as well as the individual components. As an example, a single link failure may impact the bandwidth and latency of a particular link, but is unlikely to cause a resource to become inaccessible.

Cloud service providers are deploying various forms of liquid cooling in data centers to handle the high-power dissipation of CPUs, accelerators and network switches. Google described direct-to-chip liquid cooling for their TPU 3.0 pods. Microsoft has announced two-phase immersion cooling tanks for their future data centers. The advances in cooling systems should have a significant impact on overall system reliability, but perhaps have an even larger impact on the reliability of the optical components.

The transition to CPO will allow the industry to consider new standards as everything from form factors to bandwidth density will require re-factoring interconnect in the data center. This will be an opportunity to consider new optical connectors that will bring higher reliability to one of the least reliable components in the data center. Expanded beam optics connectors, for example, could significantly improve overall reliability.

The implications of a local failure on a module, or of an entire module will need to be considered in the overall operation of the data center. The data center operator will consider the routing path diversity, slack capacity and mix of applications running on specific nodes when considering the severity of errors. Many applications can be highly tolerant of faults. Anecdotally, there are applications that only experience a 2% performance degradation when interconnect bandwidth is reduced by 50%. The data center operator will consider the overall economics of the data center operation when considering the impact of failures.

Training of AI neural networks is inherently a parallel computation. Two strategies are applied for decomposing the computation onto processing elements. These are known as model parallel and data parallel.

Model parallelism is associated with an instance of the network being evaluated on a single data sample. The nature of the parallel computation on a model result in very high local bandwidth between processing elements and/or the local memory system. Consequently, it is preferred to keep the entire model on a single package (ASIC and associated memory) where the cost of communication is low. For some neural networks, the model is too large to fit on a single package and will require several packages. As the bandwidths are still quite high, the economics drive towards colocation of the packages on a printed circuit board. This results in an on-board network which may be implemented in variety of forms. The compute module may implement network switching or just have an electrical interface to an on-PCB switch.

Data parallelism is a consequence of learning through examples (or samples). Training involving evaluating the model independently on multiple samples. For each example, adjustments to the model are proposed. These adjustments are then aggregated, filtered, and distributed across the entire training system. The aggregation is typically handled through an all-reduce communication step and comprises the bulk of data on the system-wide network. Traffic patterns are mostly uniform and all-to-all. One possible embodiment of an AI training node has multiple processing elements attached to a network switch. These are colocated on the same printed circuit assembly. The global network is implemented using a full mesh topology. The full mesh is the most economically efficient network for uniform all-to-all traffic as messages only traverse a single network link.

In architecting the system, a system designer is supplied with modular components which can be used across multiple designs. These components are not perfectly matched for a particular system so the designer must do a cost/benefit analysis on the various ratios of components, including optical links to determin the best design. The system designer often has limited choices, such as "this link wants to be 300 Gb/s. Do I use 200 Gb/s or 400 Gb/s?." A concrete example would be a system with 128 processing elements. The node design will have 4 processing elements on a PCB with some number of network switches. Each processing element has six 200-Gb/s injection ports, 1.2-Tb/s in total. The switch has 64 200-Gb/s ports. Here, the designer connects the 24 ports from the processing element to the switch. The system will have 32 nodes and so 31 switch ports are used to construct the fully meshed network. There are 9 ports left as spares.

For uniform traffic, this fully meshed network is overprovisioned and has excess bandwidth. Fully meshed networks only need 1 interswitch port for every processor port, the actual ratio is 31:24 and the network has roughly 30% residual capacity. This type of network can redistribute traffic around failed or congested links at the expense of needing to traverse more links. The example network can tolerate approximately 3 failed links per switch before the network becomes degraded.

A first order model of link reliability looks at the probability of having any 4 links fail on a switch. This is given by the binomial formula describing the overall system reliability as $\text{combin}(4,31) \times (\text{pf})^4 \times (1 - \text{pf})^{27}$. Where "pf" is the probability of failure of any single link. This calculation is analogous to the probability of getting 4 heads (link failures) when throwing 31 coins (switch ports). Assuming the probability of failure is small, we can multiply that by 32 (the number of switches) and obtain the approximate probability of network degradation. If we want the probability of the system being degraded to be less than 0.01, the associated link failure probability is 0.01 (by coincidence).

Supposing the repair policy is "links-fail-in-place" with requiring a 99% probability that the system is not degraded at 5 years, we can use the above failure probability to get the target link FIT rate. We allow 0.01 link failures every 5 years which corresponds to a per-link FIT rate of 228.

If we allow 5 links to fail on a switch, there is some loss of total available bandwidth, roughly 4%. This corresponds to a link failure probability of 0.02 or about 450 FIT for the 99% level at 5 years.

Note that most applications are not using all the available network bandwidth 100% of the time. Loss of total bandwidth turns into some decrease in application performance. Experience suggests that there is about 1% application slowdown for 5% bandwidth loss. This is for bandwidths around the desired configuration only.

Degraded operation may or may not be acceptable to the data center operator. Large private data centers tend to have minimal services to the internal customers and will focus on the most cost-effective solution, even if some jobs experience a slight slowdown. Cloud Service Providers tend to provide more stringent SLAs where even a 4% bandwidth loss is unacceptable.

9.5 Power and cooling

Power, cooling, and cost are important figures of merit for PICs. See Chapter 1 for a thorough discussion of figures of merit for all the applications of PICs in Data Centers and HPC. This section discusses power and cooling specifically for PICs in heterogenous computing systems.

The move to hyperscale data centers is exploding, driving a major shift in both global power use and power distribution within the data center. Today, hyperscale data centers are consuming roughly 2% [11] of the world's energy supply. Improvements in data center infrastructure, most importantly power delivery efficiency, have allowed the overall capacity to grow faster than the energy consumption.

While overall compute efficiency has improved, network efficiency has not, and has exposed the network as an important power consideration for overall energy consumption. Further, as Fig. 9.10 shows, each successive generation in network transceivers is increasing the power over the previous generation.

The network component has grown from 3% as shown in Fig. 9.11 from 2016 to around 10% [12] of the overall data center power:

Not only is the power in the network increasing, the power density of the network components is increasing to the point where they exceed the CPU, GPU or other accelerators. Importantly, the network switch itself has the largest thermal design point as shown in Fig. 9.12.

It is also important to point out that the energy cost (OpEx) and infrastructure cost (CapEx) of the network is considered overhead by the hyperscale data center operators. While customers will pay for higher-performance virtual machines, it is difficult to show clear customer value for network components that are necessary but do not have customer visibility for value.

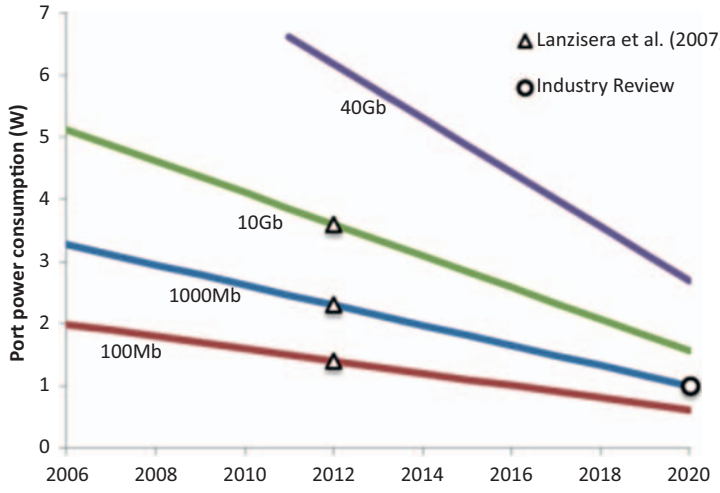

Figure 9.10 Network transceiver power vs. generation.

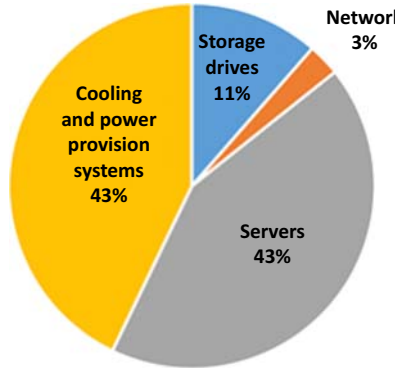

Figure 9.11 Datacenter power breakdown by component type.

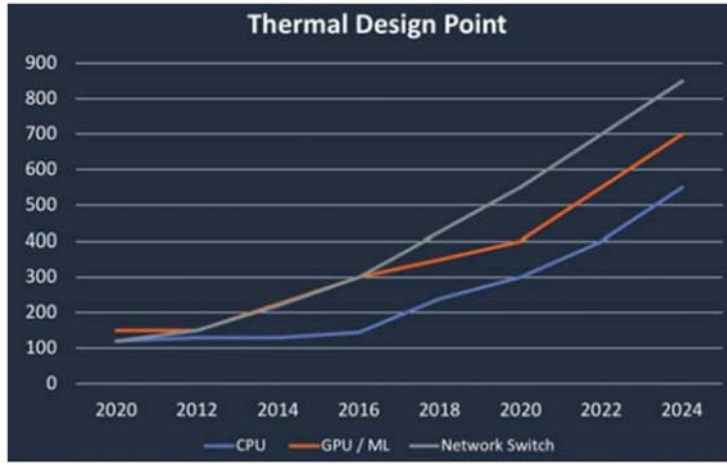

Figure 9.12 Power (Thermal Design Power) by component.

Traditionally, data centers have used air cooling where cold air is moved across hot components and the hot air cooled using an air handler. With higher power density components, different types of cooling systems are being deployed that have implications on optical components.

Liquid cooling can be in the form of cold plates that cool individual components or immersion where entire circuit boards of components are in direct contact with a fluid. Immersion can be single phase, where a liquid, such as an oil, is pumped over the surface of the hot components. Two phase immersion uses a low boiling temperature liquid to boil the liquid where it is in contact with any hot surface. The boiling process creates a passive circulation process and the liquid is recovered as it is evaporated.

9.6 Cost: current costs and path to $0.02/Gbps

The insatiable demand for higher performance data centers with a broad mix of workloads from artificial intelligence through storage are driving higher capacity interconnects for heterogeneous components. Moving optical interconnect to device endpoints is an obvious way to meet these demands, but with a higher price.

Today, optical links are about 30 × more expensive than copper links as shown in Fig. 9.13.

The benefit of optical reach (the distance between two interconnected elements), bandwidth and physical volume provide enough value to offset the cost differential. Tighter integration of optics with the migration from pluggable modules to OBO to CPO and eventually to integrated optics will dramatically reduce the cost of optical links.

While integration is key to reducing the costs of optical links, the cost will remain substantially higher than electrical interconnect without further reductions. Transforming optics from a small-scale, boutique technology to a high-volume, standardized technology will provide the second wave of cost reductions that will bring optical links to similar cost structures as electrical links. Tight integration into HVM that resemble CMOS production flows, including design kits (PDKs), manufacturing, wafer-level testing and packaging will bring a scale of economics to provide another 5−10 × in cost reduction (see Chapter 10 for more information). This HVM will also increase overall device reliability.

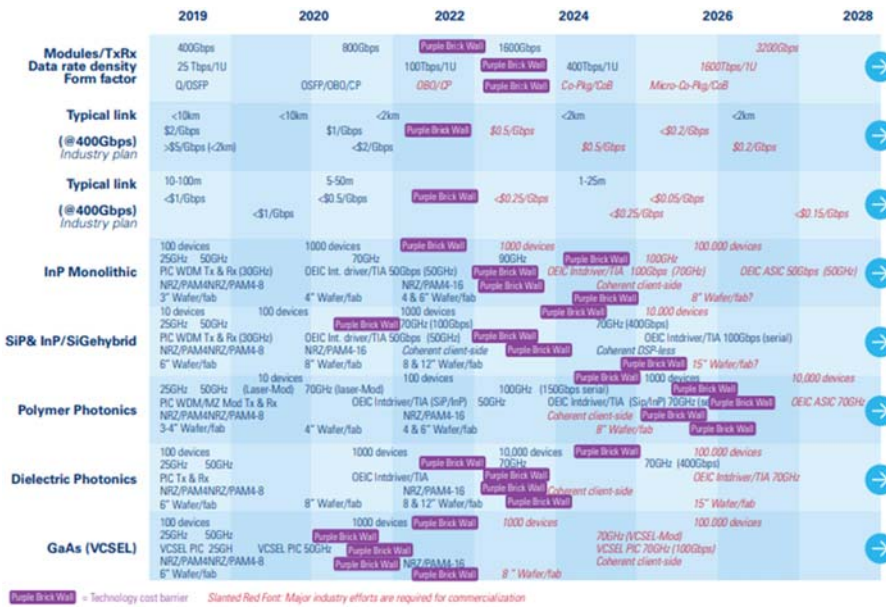

Figure 9.13 Cost of network interconnect components.

Ayar Labs has recently demonstrated an HVM flow that includes wafer level testing of optical components. Fig. 9.14 shows a wafer prober with optical interfaces performing class test on a wafer.

The wafer level testing produces test results on wafer maps that are familiar to anyone who has been associated with the production of traditional CMOS components. Fig. 9.15 shows test results for individual die locations are used to select die for packaging, avoiding costly die level tests and minimizing the chances of packaging nonfunctional die.

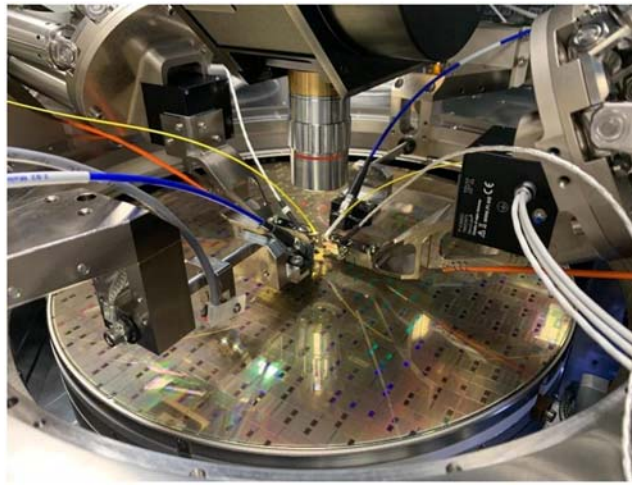

Figure 9.14 Wafer test setup for optical components.

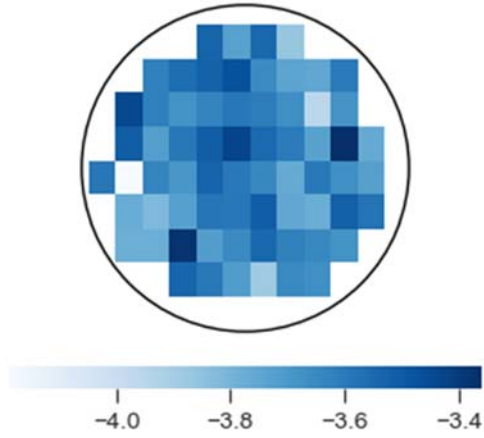

Figure 9.15 Wafer-level test results.

The move to HVM will only happen with robust market demand. The incredible growth of cloud-based data centers and the projected growth of optics for healthcare applications should clearly motivate the transition to HVM. Yole [13] has consistently projected a compounded annual growth rate (CAGR) greater than 40%, and Cisco projects that the historical optics CAGR of >20% will continue in the future. The future may see additional cost reduction by leveraging technologies that are already in very high-volume production. For example, Avicena has publicly discussed using consumer grade micro-LEDs and sensors in place of lasers and photodetectors.

9.7 Fabric architecture: control and data planes, scheduling

The previous sections compel us to consider reconfigurable photonic fabrics capable of configuring the HCS interconnection topology that would be optimally matched to the data flow pattern for the given workload, or even reconfiguring during the run time. It is extremely challenging to realize scalable and low-latency control planes for such a reconfigurable optical-circuit switching (OCS) driven by the dynamicity of the changing traffic patterns unless it is precomputed and prescheduled. The heterogeneity of the future HCS systems and disaggregation can add complexity to this process. In standard or dynamic circuit networks' physical layers, 100 μs or longer timescale reconfiguration may be sufficient, but flow-switching or burst-switching should achieve reconfigurations at below 100 μs, while packet-switched networks must achieve switching at much faster time-scale than the length of the packets (<1 ns). Seen from the applications or workload, job-level reconfigurations can be at timescales longer than 1 ms, while flow-level reconfigurations and packet-level reconfigurations should achieve <100 μs and <10 ns respectively. In terms of the control plane, depending on the scale of the HCS network and the scheduling algorithm, the centralized SDN control plane may be able to achieve and complete reconfiguration of the HCS networks at 1 ms, while faster reconfiguration should resort to distributed hardware control using FPGAs or ASICs.

The benefit of reconfigurations of communication networks in computing systems have been discussed from the perspective of efficiently and effectively utilizing the available resources (processing, memory, and communications) [14–17]. In particular, mitigating hot-spot creations in computing systems [15,18,19] can be where optical reconfigurations can prove to be very useful. The empirical studies indicate that HCSs contain hot-spot links. The hot-spot links are changing over time but can last for a long duration or persist most of the time. Losses can occur on links with low average utilization indicating that the losses are due to momentary bursts and that loss in data can happen even in under-utilized links due to bursts, and some links do not form hot-spots therefore (optical) reconfiguration could off-load the hot-spot links. These findings indicate that a static optical circuit network may not be effective for future HCSs and a dynamically reconfigurable

optical network should be considered. But the time scale of these bursts at <25-μs casts serious challenges for scheduling and for control plane designs across the data center network. Even if optical switches can reconfigure in less than 1 ns, the control plane cannot achieve coordination between all the compute nodes in the data center in such a short amount of time. In a high-performance computing system running a single threaded application with predictable changes in traffic patterns (e.g., map-reduce application transition from map-phase to reduce-phase), such a reconfiguration is conceivable if a guard time is incorporated, but it is difficult to predict such traffic patterns in a data center running many heterogeneous workloads simultaneously. Some studies are underway to apply machine-learning methods to statistically predict data flow patterns within the data centers [20,21].

The challenges in implementing centralized control planes for optical reconfiguration also raise an interesting question regarding the viability of optical-packet switching (OPS) and optical-burst switching, in addition to dynamic OCS in future data centers. For OCS based intradata center networks, we assumed the central control plane triggering reconfiguration of optical switches for dynamical reconfiguration of optical circuits. The scalability of this method depends greatly on scheduling algorithms and on the dynamicity and the load of the traffic pattern. The NEPHELE project [22] introduced WDM ring network with optical reconfiguration based on TDMA time-slot allocated by a SDN (OpenFlow) control plane to schedule these applications, and showed that the span can reach 48% when short-term load dynamicity is high. Such centralized schedulers inevitably add scheduling delays which can become unacceptably high in large data center networks. Just to poll the traffic demands, the delay T_D scales as $O(N^2)$, or $T_D = B \times N^2/C_{BW}$ [23], which corresponds to 4 ms for $N = 5000$ racks, $B = 16$ bits of scheduling bits ignoring framing bits, $C_{BW} = 100$ Gb/s. Computation time to calculate the optimal topology for reconfiguration also scales as $O(N^{1.5})$ for Hopcroft Karp [24], $O(N^2\log^2(N))$ for Solstice, and much higher order for Eclipse.

In OPS networks, the control plane can be completely distributed, and it can handle short burst traffic on a packet-by-packet basis to support scalability. While lack of optical buffers has been a main challenge, contention resolution schemes in wavelength-time-space domains have demonstrated to be effective [25]. Recent publications [26] also noted the value of including SDN in OPS networks.

A number of important demonstrations and studies have taken place on this topic. Microsoft Research has demonstrated a SIRIUS switch network [27] based on tunable lasers and AWGR wavelength routing arranged in the Thin-CLOS [28] architecture different from the all-to-all arbitration-free network of LION discussed in [29]. Hence, due to the requirement of optical contention resolution at the receiving node, scheduling of traffic can become even more challenging for SIRIUS. Ballani [27] correctly points at the daunting challenges of on-demand scheduling, by sending the data center demand matrix measured at the node, to a scheduler that calculates and assigns communication timeslots to the nodes. While such an approach may be viable for relatively slower optical switches at coarse timescales of ∼millisecond, it is not efficient and practical for fast switching (microsecond and below) in large scale networks. Instead, Sirius adopts a scheduler-less design [30] as an extension of Valiant load balancing

through detouring to all other intermediate nodes, which then forward the traffic to its destination node to handle nonuniform traffic demands, while keeping the control plane to be distributed. Another notable recent demonstration is UCL's PULSE [31] scheduler which achieves sub microsecond scale scheduling for star-coupler based sub-networks that are disjoint by using multiple transceivers with wavelength-timeslot selective architectures for each disaggregated node in 64 node racks. At the reconfiguration circuit duration (*epoch* = 120 ns), the scheduling algorithm is shown to achieve up to 93% throughput with disaggregated nodes.

9.8 Photonic-integrated circuits for heterogeneous computing systems

Heterogeneous computing is pressuring bandwidth, reach, and latency requirements, and growth in high value application workloads, such as machine learning model training, are outpacing improvements in current technology and product roadmaps. As illustrated in Fig. 9.16 [Ref. DARPA PIPES Proposers' Day presentation slides Gordon Keeler], traditional approaches of placing discrete photonic and electronic components are not scalable nor power efficient. Fig. 9.17 illustrates the general trend of packaging to cointegrate PICs and electronic-integrated circuits, to realize disaggregated HCS in a scalable and power-efficient manner.

A new generation of PICs can address these opportunities and challenges in the form of (1) photonic *I/O*s (transceivers and receivers closely integrated electronic ICs forming compute nodes) and (2) photonic switches with all-optical switching, hybrid photonic-electronic switching, or electronic switching with photonic *I/O*s. This Section 9.8 of the Chapter describe both types of PICs for HCS.

9.8.1 Photonic-integrated circuits for photonic I/Os

One example of PICs for photonic *I/O* is an early prototype of the Ayar Labs TeraPHY chiplet [10], shown in Fig. 9.18. The PIC integrates electronics (digital

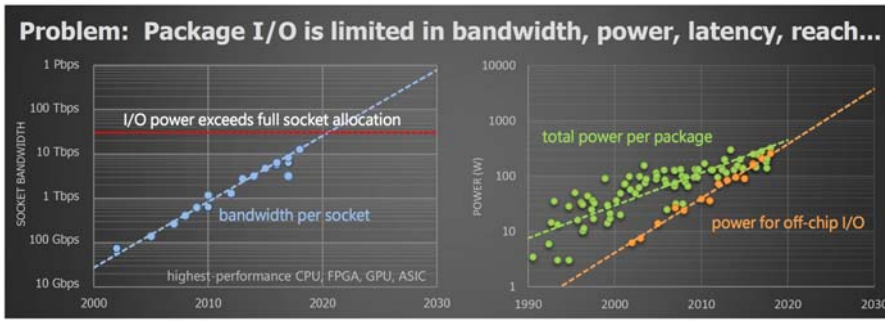

Figure 9.16 *I/O* bandwidth scaling and power-problems of traditional packaging methods.

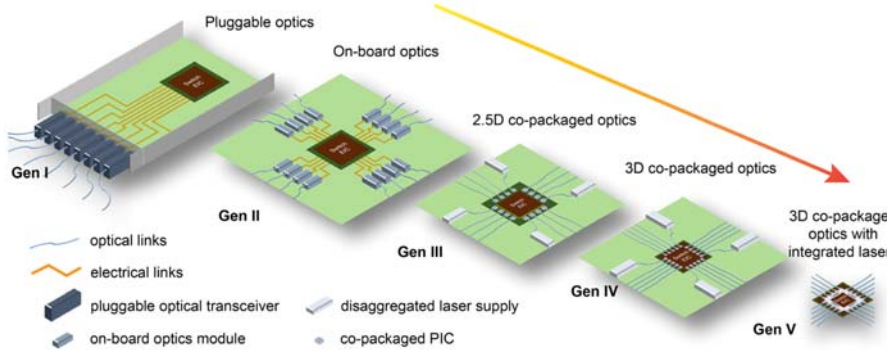

Figure 9.17 Generations of optics and the evolution of copackaging technologies used in data center applications. The generational progression drives tighter integration between network switching and optical *I/O* that will probably culminate with 3D copackaged optics with integrated lasers on chip.
Source: From Appl. Phys. Lett. 118 (2021) 220501. https://doi.org/10.1063/5.0050117.

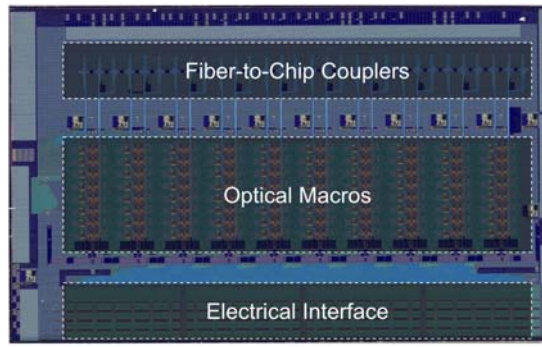

Figure 9.18 Prototype TeraPHY PIC from Ayar Labs targeting multi-Tbps chip-to-chip communications.

circuits, high-speed analog-mixed signal circuits) with photonics (fiber couplers, waveguides, couplers, filters, optical modulators, photodetectors, etc.) resulting in a single chip with $\sim$50 million electrical transistors and $>$1000 optical devices. This chiplet achieves a bandwidth density of $\sim$0.25 Tbps/mm with a clear path to $>$1 Tbps/mm using the same architecture.

The PIC contains three main architectural regions: an electrical interface, optical macros, and fiber-to-chip couplers. The electrical interface connects to a host system-on-chip (SoC) interface to transmit and receive electrical data. The optical macros serialize/deserialize the electrical signals to transmit and receive at a higher data rate per optical channel. In this example, the optical macros utilize silicon

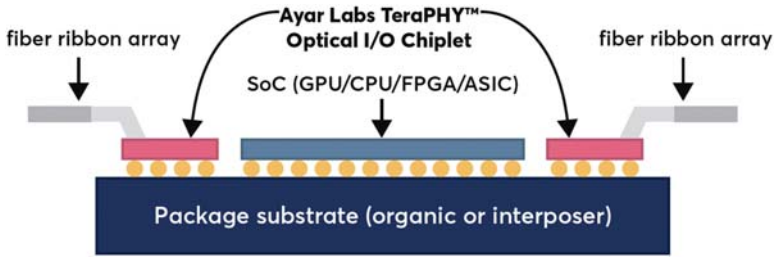

Figure 9.19 A schematic example of flip-Chip attached packaging of TeraPHY optical *I/O* chiplets with electronic ICs.
Source: Reproduced with permission from AyarLabs Technical Brief, Optical I/O Chiplets Eliminate Bottlenecks to Unleash Innovation https://ayarlabs.com/technical-brief-optical-i-o-chiplets-eliminate-bottlenecks-to-unleash-innovation/#unlock.

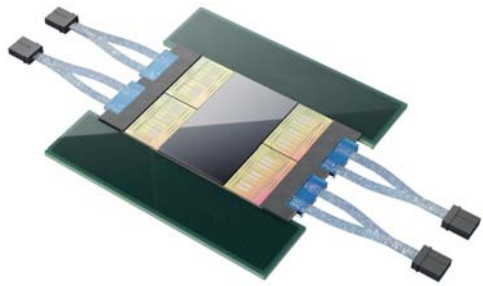

Figure 9.20 TeraPHY optical *I/O* multichip package integration example from Ayar Labs.
Source: Reproduced with permission from AyarLabs Technical Brief, Optical I/O Chiplets Eliminate Bottlenecks to Unleash Innovation https://ayarlabs.com/technical-brief-optical-i-o-chiplets-eliminate-bottlenecks-to-unleash-innovation/#unlock.

microring technology to implement wavelength division multiplexing (WDM) to maximize the full-duplex bandwidth per fiber. Overall, the TeraPHY chiplet leverages features only recently available through photonics foundry and advanced packaging and offers a promising direction to solve critical challenges facing HCS.

As discussed earlier, optical I/O PICs, such as TeraPHY prototypes, need to be closely packaged together with electronic ICs that generate and receive data to reduce power consumption and benefit from high bandwidth density. Figs. 9.19 and 9.20 illustrate an example of such packaging for TeraPHY Optical I/O PIC chiplets and electronic ICs (e.g., ASICs, GPUs, CPUs).

Such heterogeneous compute nodes with integrated PIC *I/O*s as shown in Fig. 9.20 can readily interconnected by another PIC with wavelength routing

capabilities as shown in Fig. 9.21A. In this case an arrayed waveguide-grating router (AWGR) with NxN cyclic wavelength routing capabilities can achieve all-to-all interconnection between the *N* heterogeneous compute nodes. Fig. 9.21B shows a fabricated silicon photonic AWGR with integrated silicon photonic transceivers, and Fig. 9.21C illustrates a schematic of a compute cluster formed by *N* heterogeneous compute nodes and a silicon photonic AWGR PIC.

9.8.2 Photonic-integrated circuits for enabling disaggregated reconfigurable heterogeneous computing systems

In Sections 9.1 and 9.2, we described the benefit of combining heterogeneous system together with reconfiguration so that application aware configuration of network interconnect topologies matching the data flow pattern driven by the application. The energy efficiency and throughput enhancement (often 50 or more [2,3]) achievable in domain-specific computing relative to the general purpose computing can be realized in a large systems consisting of heterogeneous domain-specific computing elements configured for optimal execution of a given workload. To achieve such benefits for various applications, optical reconfiguration of heterogeneous computing resources is desirable. Such optical reconfigurations are achievable by optical switches described in Chapter 8 of this Monograph and recent publications [32]. Commercially available optical switches are typically electronic switches (such as Tomahawk4 switches [33]) as shown in Fig. 9.22. This Humboldt System corresponds to Gen I in Fig. 9.17 utilizing pluggable QSFP optical transceivers. Gen III version may conceptually look like Fig. 9.23 utilizing the optical *I/O* PICs shown as "PE" in the figure exemplified as AyarLab's TeraPHY in Fig. 9.17.

On the other hand, an all-optical-switch PIC can achieve the desired reconfiguration of the network topology while eliminating the needs for repeated optical−electrical−optical conversions. One such example is Flex-LIONS (Flexible Low-latency interconnect network switch) [34,35]. Fig. 9.24 illustrates the utility of such a reconfigurable wavelength routing switch with all-to-all interconnect capability showing hot-links can be assigned additional wavelengths to achieve higher throughput and energy efficiency (note that typical transceivers and compute nodes consume similar amounts of power even when there are no actual data movements therein). Fig. 9.25B shows the silicon photonic Flex-LIONS PIC consisting of the arrayed waveguide grating router, microresonator ring (MRR) add/drop filters, and an MRR crossbar switch monolithically integrated on a single silicon-on-insulator platform. Fig. 9.25A shows the packaged Flex-LIONS. Fig. 9.25C shows all-to-all interconnect topology offered by the Flex-LIONS before reconfiguration and Fig. 9.25D illustrates an example of bandwidth-steered interconnect topology offered by the Flex-LIONS after a reconfiguration (Figs. 9.24 and 9.25).

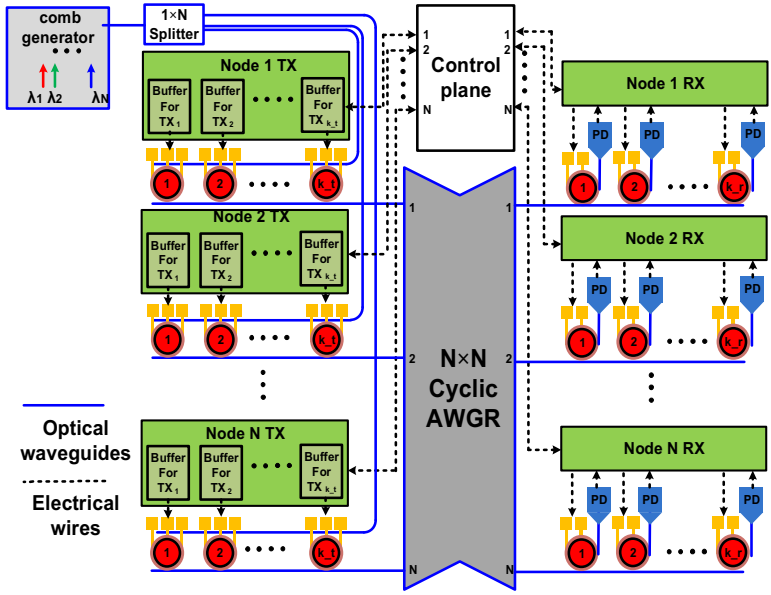

(A)

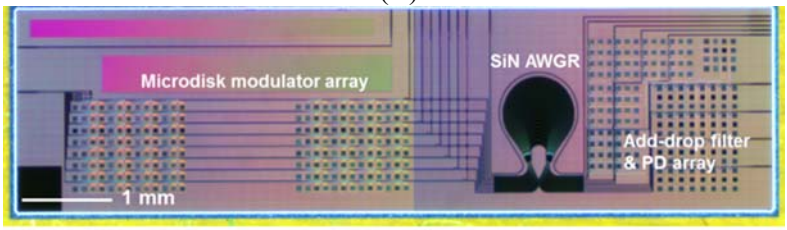

(B)

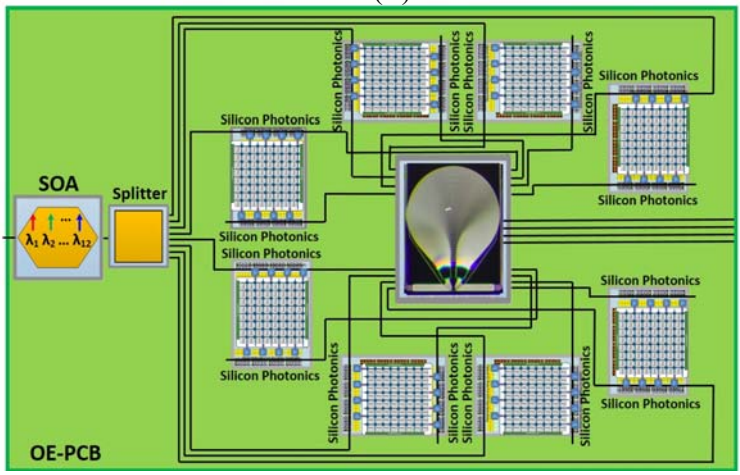

(C)

Figure 9.21 (A) Low-latency interconnect network switch architecture involving a cyclic $N \times N$ arrayed waveguide-grating router, compute nodes with embedded electronic switches

(*Continued*)

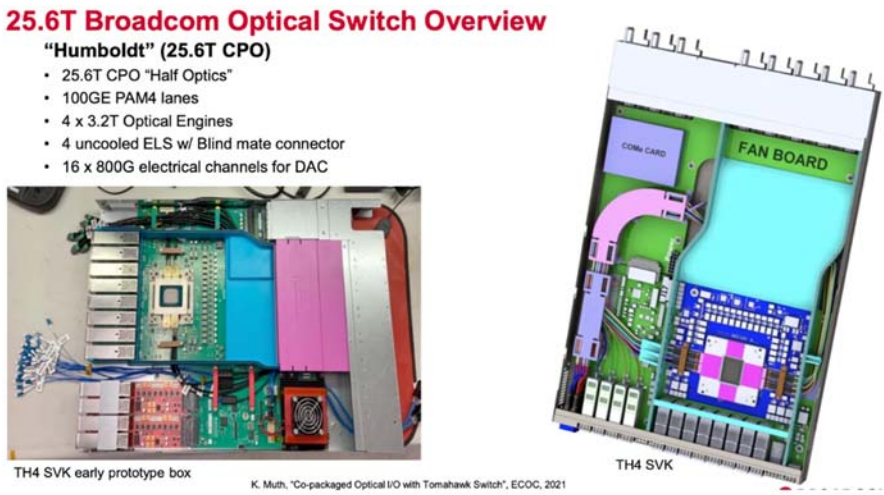
25.6T Broadcom Optical Switch Overview
"Humboldt" (25.6T CPO)
• 25.6T CPO "Half Optics"
• 100GE PAM4 lanes
• 4 x 3.2T Optical Engines
• 4 uncooled ELS w/ Blind mate connector
• 16 x 800G electrical channels for DAC
TH4 SVK early prototype box
K. Muth, "Co-packaged Optical I/O with Tomahawk Switch", ECOC, 2021
COMe CARD
FAN BOARD
TH4 SVK

Figure 9.22 25.6-Tb/s broadcom optical switch "HUMBOLDT" system 64 × QSFP-DD/OSFP 25.6 Tbps with Tomahawk4 in 2RU.
Source: From K. Muth "Co-packaged Optical I/O with Tomahawk Switch," ECOC 2021.

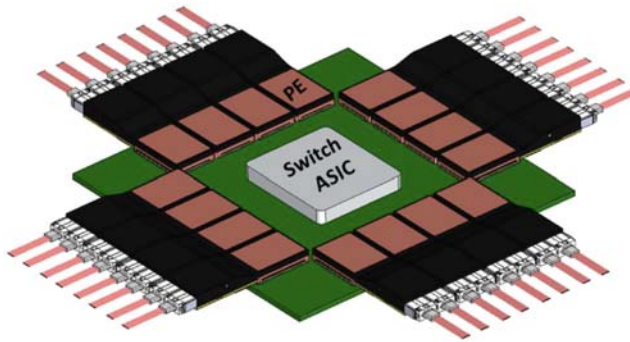

Figure 9.23 A conceptual schematic of a copackaged-optics switch subassembly.
Source: From S. Fathololoumi, D. Hui, S. Jadhav, J. Chen, K. Nguyen, M. Sakib, et al., 1.6 Tbps silicon photonics integrated circuit and 800 Gbps photonic engine for switch co-packaging demonstration, J. Light. Technol. (2020). https://doi.org/10.1109/JLT.2020.3039218.

and wavelength division multiplexing silicon photonic transceivers, and an optical frequency comb (wavelength division multiplexing) laser. (B) Fabricated SiPh low-latency interconnect network switch and (C) board level integration using $N \times N$ arrayed waveguide-grating router for p number of multisocket compute nodes with silicon photonics transceivers with μ number of external *I/O* ports (shown is for $p = 8$, $\mu = 4$, and $N = p + \mu = 12$).
Source: Courtesy of S.J. Yoo, Optical-label based packet routing system with contention resolution in wavelength, time, and space domains, in: Optical Fiber Communications Conference (OFC), Anaheim, 2002.

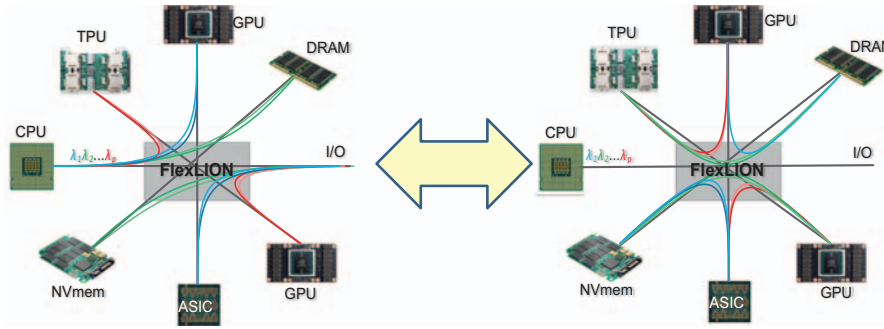

Figure 9.24 Optical reconfiguration of networking topologies interconnecting heterogeneous computing resources can significantly enhance energy efficiency and throughput of the system adapting to the required data movements driven by the workload. Shown is an example of FlexLION which achieves reconfigurable wavelength routing between heterogeneous compute nodes.

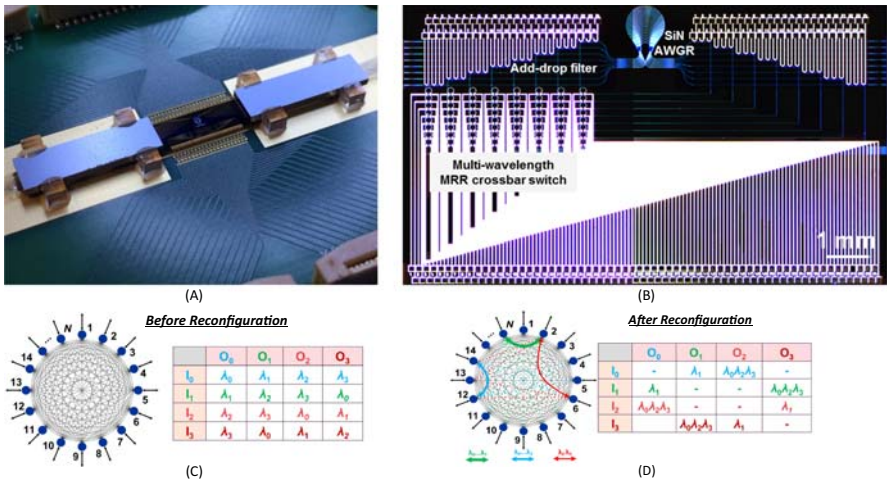

Figure 9.25 (A) Photograph of the integrated and packaged Flex-LIONS module. (B) Microscope image of the fabricated 8×8 SiPh Flex-LIONS ($N = 8$, $b = 3$) chip. (C) All-to-all interconnect topology offered by the Flex-LIONS before reconfiguration. (D) An example of bandwidth-steered interconnect topology offered by the Flex-LIONS after a reconfiguration.

9.9 Conclusion

Driven by ubiquitous adoption of machine learning applications in hyperscale data centers, the exponential growth of data ingestions and processing in data centers have rapidly accelerated to doubling every 3.5 months. Scalability, energy efficiency, throughput, and disaggregation are all very important in future computing and data

systems. Heterogeneity of computing resources is important because application-specific computing systems optimally designed and configured for a given workload offer much higher (often 50 or more) [2,3] energy efficiency and throughput compared to general-purpose computing systems. Optical interconnects of such heterogeneous computing resources allow disaggregation of the resources while offering scalability and energy efficiency independently of distance. PICs closely integrated with electronic ICs (EICs) such as heterogeneous computing elements can form compute nodes with heterogeneous domain-specific accelerators. PICs with passive or active reconfiguration capabilities can interconnect such heterogeneous computing resources to bring locality to the data in a system handling sparse and nonuniform data transactions. In multitenant data centers, such wavelength reconfigurations also allow partitioning of heterogeneous compute resources.

Data centers are evolving from moderately homogeneous server configurations providing first and third-party services that were well adapted to the compute and network topologies of the time. New demands from machine learning training and inference as well as diverse applications from third parties are populating data centers with highly heterogenous processing architectures. High throughput GPUs used for machine learning, visualization and hosting along with custom and third-party accelerators are creating diversity of hardware that are creating unique demands for interconnect fabric.

This chapter has covered the problem that has emerged with heterogenous systems along with the specific application classes that are driving the changes in requirements for data center fabrics. These applications drive requirements for PICs that are different from coherent data center interconnects (Chapter 6) and PICs for switching applications (Chapter 7). The detailed PIC design and resulting challenges are detailed for heterogenous systems.

References

[1] W.A. Wulf, Hitting the Memory Wall, ACM SIGARCH Computer Architecture News, 1996.
[2] W.J. Dally, Efficient embedded computing, Computer 41 (2008) 27–32.
[3] A. Abnous, Ultra-low-power domain-specific media processors, VLSI Signal. Process. (1996) 461–470.
[4] A. Naeemi, Optical and electrical interconnect partition length, IEEE Photonics Technol. (2004) 1221.
[5] R. Beausoleil, Nanoelectronic and nanophotonic interconnect, in: Proceedings IEEE, 2008, 230–247.
[6] G. Liu, Architecture and performance studies of 3D-Hyper-FleX-LION, in: SC20 Proceedings, 2020, 1–16.
[7] IEEE, International Roadmap for Devices and Systems™ 2020 Edition, 2020. <https://irds.ieee.org/editions/2020>.
[8] 100G Lambda, 100G Lambda MSA announces release of new specifications, 2020. <https://100glambda.com/>.
[9] M. Sysak, CW-WDM MSA technical specifications. <https://cw-wdm.org/?wpdmdl = 2092>, 2021.

[10] M. Wade, An error-free 1 Tbps WDM optical I/O Chiplet, in: Optical Fiber Communication Conference (OFC), Optica Publishing, 2021, p. F3C.6.
[11] C. Garcia, The real amount of energy a datacenter uses. Retrieved from AKCP. <https://www.akcp.com/articles/the-real-amount-of-energy-a-data-center-use/>, 2022.
[12] E. Innovation, Energy innovation. Retrieved from Energy Innovation. <https://energyinnovation.org/2020/03/17/how-much-energy-do-data-centers-really-use/>, 2020.
[13] Yole, 2020−2026 optical transceiver growth forecast. Retrieved from Yole. <http://www.yole.fr/iso_album/illus_optical_transceivers_2020_2026_revenuegrowthforecast_yole_july2021.jpg>, 2021, July 21.
[14] K. Chen, An optical switching architecture for data center networks with unprecedented flexibility, IEEE ACM Trans. Netw. 22 (2014) 498−511.
[15] G. Porter, Integrating microsecond circuit switching into the data center, SIGCOMM (2013) 447−458.
[16] R. Tessier, Reconfigurable computing architectures, Proc. IEEE 103 (2015) 332−354.
[17] G. Wang, Programming your network at run-time, in: Proc. First Work. Hot Top. Softw. Defin. Netw., 2012, 103−108.
[18] T. Benson, Network traffic characteristics of data centers in the wild, in: Proc. 10th ACM SIGCOMM Conf. Internet Meas, 2010a, pp. 267−280.
[19] T. Benson, Understanding data center traffic characteristics, SIGCOMM Comput. Commun. (2010) 92−99.
[20] X. Chen, Machine-learing-aided cognitive reconfiguration, J. Opt. Commun. Netw. (2020).
[21] M. Balanici, Machine learning-based traffic prediction for optical switching resource allocation in hybrid intra-data center networks, in: Optical Fiber Communications Conference and Exhibition (OFC), 2019.
[22] Bakopoulos, P. (2018). NEPHELE: an end-to-end scalable and dynamically reconfigurable optical architecture for application-aware SDN cloud data centers. IEEE Commun. Mag. *56*, (178−188).
[23] E. Zahavi, ODCNs architectures fundamental limits, in: Opt. Fiber Commun. Conf, 2019.
[24] J.E. Hopcroft, An $n^{5/2}$ algorithm for maximum matchings in bipartite graphs, SIAM J. Comput. 2 (4) (1973).
[25] S.J. Yoo, Optical-label based packet routing system with contention resolution in wavelength, time, and space domains, in: Optical Fiber Communications Conference (OFC), Anaheim, 2002.
[26] W. Miao, SDN-enabled OPS with QoS guarantee for reconfigurable virtual data center networks, IEEE OSA J. Opt. Commun. Netw. 7 (2015) 634−643.
[27] H. Ballani, Sirius: a flat datacenter network with nanosecond optical switching, in: ACM SIG on Data Comm SIGCOMM, 2020, pp. 782−797.
[28] Y. Yin, AWGR-based all-to-all optical interconnects, 2022.
[29] R. Proietti, A scalable, low-latency, high-throughput, optical interconnect architecture based on arrayed waveguide grating routers, J. Light. Technol. 33 (2015) 911−920.
[30] C.S. Chang, Load balanced Birkhoff-von Neumann switches, part I: one-stage buffering, Comput. Commun. (2002).
[31] J.L. Benjamin, Scaling PULSE data center network architecture and scheduling optical circuits in sub-microseconds, in: Optical Fiber Communication Conference (OFC), 2020, p. W1F.3.
[32] S.J. Yoo, Prospects and challenges of photonic switching in data centers and computing systems, J. Light. Technol. 40 (2022) 2214−2243.

[33] T.P. Morgan, Broadcom launches another tomahawk into the datacenter. The next platform. <https://www.nextplatform.com/2019/12/12/broadcom-launches-another-tomahawk-into-the-datacenter/>, 2019.
[34] X. Xiao, Silicon photonic Flex-LIONS for bandwidth-reconfigurable optical interconnects, IEEE J. Sel. Top. Quantum Electron. 26 (2) (2020) 1−10.
[35] X. Xiao, Multi-FSR silicon photonic flex-lions module for bandwidth-reconfigurable all-to-all optical interconnects, J. Light. Technol. 38 (12) (2020) 3200−3208.

This page intentionally left blank

Photonic integrated circuit design methods and tools

10

Twan Korthorst[1], *Wim Bogaerts*[2], *Duane Boning*[3], *Mitch Heins*[1] *and Barton Bergman*[4]

[1]Synopsys, Inc., Mountain View, CA, United States, [2]Ghent University - imec, Ghent, Belgium, [3]Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, Cambridge, MA, United States, [4]Dream Photonics, Vancouver, BC, Canada

10.1 Introduction

To build silicon photonics transceivers, or photonic chips in general, one needs to go through a process of design, fabrication, and test. This chapter mostly deals with the first step in this flow but also discusses how it cannot be considered independent from the subsequent steps. A good design process captures the essence of the intended functionality and translates this into a chip design that can be fabricated and that eventually works, even in the presence of variations in the fabrication process. When we look at the design of a high data rate optical link from a system perspective, we can identify many different elements that need to be integrated: digital input signals need to be multiplexed electronically and modulated onto an optical carrier wave, which is transported over the optical fiber. At the receiver side these signals are detected by a photodetector and converted into digital signals and demultiplexed. Links can be based on simple intensity modulation (two-level or multilevel formats), or coherent modulation using both phase and amplitude to encode data. Using wavelength division multiplexing (WDM), multiple of these signals can be multiplexed on a single fiber, and it is also possible to parallelize into multiple fibers or multiple spatial modes in the same fiber.

Such a link consists of photonic and electronic chips that must work closely together and need to be integrated in a functional package with optical, electrical, and high-speed connections, and integrated thermal management. The form factor can adhere to standard pluggable modules or scale down to on-board optics, near or co-packaged optics, or full-custom modules.

In this chapter, we focus particularly on the methods for designing photonic chips, photonic circuits composed of devices as illustrated in Fig. 10.1. Like electronic chips, photonic chips are fabricated using planar fabrication processes including lithography, etching, deposition, and polishing. It is therefore logical that the design flow for photonic chips follows practices that have already been established for electronic circuit design. For instance, the fabricated geometries are described as a set of patterns (masks) for each process step (Chapter 11), building up the complete photonic chip layer by layer.

Integrated Photonics for Data Communication Applications. DOI: https://doi.org/10.1016/B978-0-323-91224-2.00008-4
© 2023 Elsevier Inc. All rights reserved.

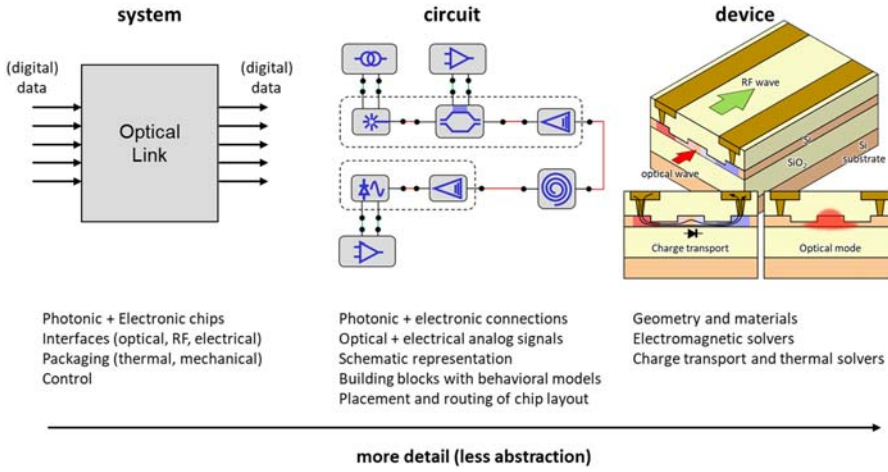

Figure 10.1 Levels of abstraction in photonic design. The high-level specifications, which dictate the desired inputs and outputs, can be used to decompose the system into circuits of functional building blocks described by efficient behavioral models. These building blocks, or devices, can be designed by simulating the physical processes within the geometry.

However, there are fundamental differences between photonics and electronics, and this also translates into the design flow [1−3]. While both photonics and electronics can be abstracted as circuits, their behavioral description is very different. Where electronic circuits consist of wiring described by Kirchhoff's conservation laws, photonic circuits consist of waveguides that can carry distinct signal waves in both directions. Moreover, these signals are modulated on a high optical carrier frequency (or wavelength) in both amplitude and phase, and multiple of these modulated frequencies can be present in a single waveguide [4].

The signal representation depends very much on the type of circuit or the application for which the chip is being designed. For instance, for the design of a simple transceiver using an on−off keying modulation scheme, it might be sufficient to encode the optical signal as an intensity. However, when using coherent communication schemes, such as DPSK or QAM16, both the intensity and the phase information is needed to accurately describe the optical signals. And when we use a WDM scheme with multiple carrier frequencies, the modulated intensities and phases for each carrier wave are needed during the simulation to accurately capture the crosstalk effects in the circuits.

Because they are described by traveling waves, photonic circuits are more reminiscent of RF circuitry, but they can scale to much higher complexity because the optical wavelengths are so much smaller. Because of that growing complexity, photonic chips cannot be designed as a monolithic element: the electromagnetic simulation of the optical waves and their interactions throughout the entire chip would become an intractable problem. Thus, a circuit abstraction is needed, where photonic building blocks are wired together, and propagating signals describe the

electromagnetic waves in the connecting waveguides. Dedicated photonic circuit design tools help the designer to decompose the circuit into hierarchical subcircuits, down to the individual building blocks.

The building blocks, photonic active and passive devices, themselves are designed using electromagnetic simulators that solve Maxwell's equations inside the geometry, modeling how the light propagates through the structure. This results in a description of how the device's waveguide ports couple to each other, and this can be encapsulated in a compact model that is then used for circuit simulation. The 3D device geometry is represented by a set of 2D geometries on different layers suitable for fabrication.

In general, device design and circuit design are kept separate. The device engineer is responsible for designing an optimized device for a specific function, and then provides the abstraction to the circuit designer who can cascade multiple devices into circuits [2]. The interface between these two worlds is a process design kit (PDK), which is described further in Section 10.2.5. It is an ensemble of information that a circuit designer needs to generate a circuit design that can be fabricated by a foundry. A PDK contains basic descriptions about the fabrication process, and a library of building blocks with both a layout representation and a behavioral circuit model. The PDK is usually provided by the foundry but is often augmented with additional optimized building blocks from in-house or third-party device designers. It is this separation of concerns that enables an ecosystem of foundries supporting a large number of fabless chip companies [5].

The abstraction at the circuit level makes it possible to define much larger and more complex photonic systems, but this comes at the cost of some accuracy: the modeled behavior of the circuit is only as accurate as the individual abstract models of the building blocks. Any effects that are not described by those models will eventually not be captured in the simulation. One aspect that is particularly critical with silicon photonic circuit design is the impact of fabrication variations on the overall performance of the circuits. With its high refractive index contrast, silicon photonic waveguides and devices are very sensitive to even minute, nanometer-scale deviations in the geometry and material distribution. In actual fabrication processes, deviations on this scale are unavoidable (e.g., changes in waveguide line width or layer thickness), and they will impact the performance of the circuit in a material way. Photonic design tools are evolving to incorporate the effects of fabrication variations into the circuit design flow to enable the designers to predict the yield of their circuit after fabrications, and then use this information to optimize the design for manufacturability (DFM) [6,7].

In this chapter, we will discuss the different aspects of photonic design. We start in Section 10.2 with the aspects of physical design of individual devices and building blocks. In Section 10.3, we make abstractions of these devices and discuss the circuit design flow. Section 10.4 briefly discusses how photonic circuit design interfaces with other elements of the system, such as high-speed electrical connections and packaging. Finally, in Section 10.5, we take a deeper look into the effects of variability and how we can design for better yield. Finally, we look ahead in Section 10.6 to the future needs for a photonic design flow.

10.2 Photonic device design

Photonic circuits transport information in the form of light, which is simply an electromagnetic wave at high frequencies (180−240 THz for the commonly used wavelengths between 1250 and 1600nm). The design of a photonic chip effectively boils down to creating a geometry that controls the propagation of this electromagnetic wave in the desired way by engineering the distribution of materials with different refractive index and absorption, and the use of electro-optic and thermo-optic effects. As we will see, the rigorous simulations that are needed for such a design are compute intensive, and do not scale well to very large areas or volumes. Therefore, to keep the design manageable, a photonic chip is broken down into smaller building blocks or devices that are individually designed, and then later combined into the circuit. This section deals with the different modeling techniques that are used for device design.

When designing the building blocks for photonic integrated circuits (PICs), the purpose is to optimize the distribution of materials for a particular optical function, and this is closely related with the fabrication processes being used. Because most semiconductor processes are planar processes (deposition of layers, and then defining a pattern into those layers), we can identify two types of problems that often require different types of design and simulation tools:

- Defining the cross-section of the devices and the chip, mainly determined by the fabrication process.
- Defining the top view, mainly determined by the mask layout of the devices and the interconnects.

Device design often requires a combination of both, while the larger-scale circuit design (see Section 10.3) can be fully handled in the top view. Cross-section tools are closely tied to the physics of the materials and devices and focus on an accurate representation of device behavior. Combining these cross-sections into a 3D geometry then makes it possible to simulate the electromagnetic propagation in a device. From this we can extract models and layout parameters that will feed into the top-view circuit design tools where the building blocks are connected.

10.2.1 Technology computer-aided design

To be able to manufacture photonic devices that together form a photonic circuit, a fabrication process is required. As discussed in Chapters 2−5, photonic devices and circuits are constructed from semiconductor materials like silicon, III−V materials like indium phosphide (InP), and dielectrics. To build the required structure, a series of thin film fabrication steps must be performed. The design process for these structures can be accelerated through process simulation tools. These tools are called process TCAD (technology computer-aided design) tools and are instrumental for the development of advanced semiconductor processes as required for integrated photonics. Fig. 10.2 shows some examples of complex geometries and dopant levels simulated with a TCAD solution.

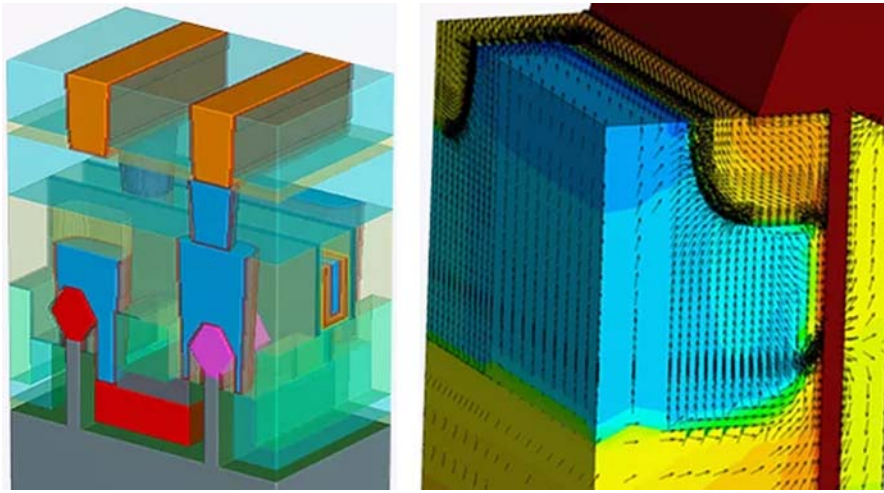

Figure 10.2 Example of 3D semiconductor process and device simulations.

TCAD solutions for electronics focus on the design, simulation, and analysis of semiconductor devices, like transistors or memory cells, modeling effects, such as carrier transport and heat transfer. Such tools are often based on finite element or finite difference solvers. The photonic equivalents are sometimes called Photonic TCAD of photonic device tools, and these also incorporate an electromagnetic solver for Maxwell's equations.

As described extensively in Chapters 2–5 a wide variety of photonic devices exist. Simply said, there are passive devices that only have optical inputs and outputs, and active devices that have both optical as well as electrical inputs and outputs. Given the fact that the devices often have features with sizes smaller than the wavelength of the light used (e.g., 1310 nm for datacom applications) optical simulation tools based on raytracing techniques [8] cannot be used. Instead, to model the propagation of light one needs to use electromagnetic (EM) or Maxwell's equations solvers. Since the early 1990s EM solvers for photonics have been developed and commercialized, and are today widely available and commonly used by scientists, researchers, and product development engineers. We can separate these solvers into mode solvers or propagation solvers.

10.2.2 Mode solvers

A photonic device is a 3D geometrical structure with material properties that determine how light interacts with the device. In many cases, light will have a preferential propagation axis. The simplest example of this is a photonic waveguide. Like an RF waveguide, the photonic waveguide is a structure guiding electromagnetic waves along one direction. The cross-section of the waveguide greatly determines the performance of the complete circuit, as propagation losses play an important role in device design. To optimize the cross-section of a photonic waveguide, mode

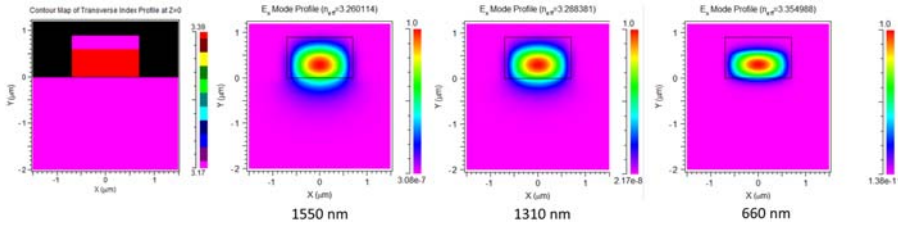

Figure 10.3 Mode profiles at various wavelengths for a straight strip waveguide.

solvers are used. A mode solver is a tool that solves Maxwell's equations on a cross-sectional mesh of the waveguide and calculates the eigenmodes of that cross-section. Each eigenmode represents a propagating wave in the waveguide with its own particular mode profile, a propagation velocity described by the effective (refractive) index, and loss. While propagating through a straight waveguide, the eigenmode does not change (though it can be attenuated) or mix with other eigenmodes. Playing with width and height and the materials that form the waveguide core and cladding, one can optimize the cross-section for a given number of guided modes, and minimize the losses by observing how much energy is contained within the structure. Fig. 10.3 shows a few examples of mode profiles for a straight strip waveguide at different wavelengths. Depending on the geometrical complexity of the cross-section (rectangle, slab, multislab) and materials used (metals are particularly difficult as they induce a very large refractive index difference and high loss), different implementations of the Maxwell equation solver provide the required accuracy at a reasonable compute cost (time, memory). There are 1D and 2D mode solvers, using predominantly effective index, finite element, finite difference, and film mode matching methods [9].

It is important to realize that curved or bending waveguides with the same cross-section as a straight waveguide still behave differently. There is a direct relation between the bending radius of the curve and energy loss. In a waveguide bend, the modal profile of the propagating light is shifting to the outer side of the curve and energy starts to radiate out (as can be seen in Fig. 10.4). As a result, photonic waveguides require careful design of the bends to provide circuit designers the ability to optimize the connections between the various devices in the circuit. Often several different optimized waveguide cross-sections and layouts are available to minimize losses when constructing a complete circuit.

10.2.3 Propagation solvers

Mode solvers calculate the distribution of light in a cross-section, but they do not tell us how light propagates when that cross-section changes due to the top-down layout that we define to obtain a 3D device geometry on a chip. To simulate this propagation, additional techniques are needed [10].

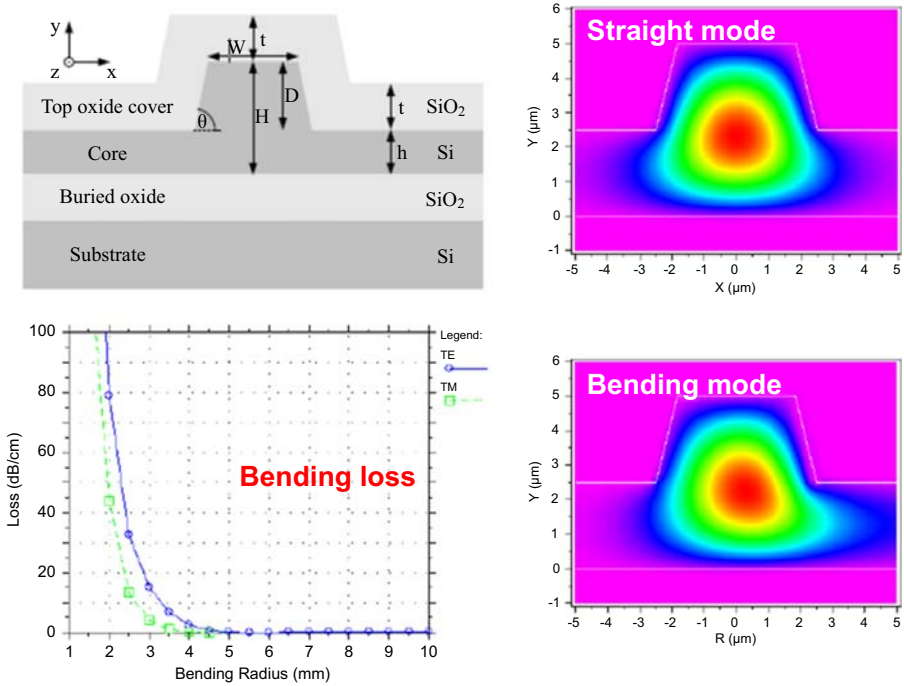

Figure 10.4 Mode profiles for straight and bend waveguides, resulting in loss for smaller bending radii because of energy leakage out of the waveguide.

10.2.3.1 Eigenmode expansion and propagation

The eigenmode expansion (EME) or bidirectional eigenmode propagation (BEP) technique builds on the mode solvers by slicing up a design into piecewise constant sections and then expressing the propagating light into a basis set of local eigenmodes for each cross-section, rigorously solving Maxwell's equations [11]. This method is intrinsically bidirectional and therefore can handle reflections. The technique is especially efficient for repeating structures like gratings, as mode calculations can be reused between identical sections. These tools can work both in 2D and full-vectorial (3D) mode, with the former being much faster and the latter being more accurate, especially in geometries with a high contrast in refractive index. While rigorously correct, the discretization into slices introduces a preferential axis, and the technique is therefore mostly useful for device geometries with a dominant propagation direction.

10.2.3.2 Beam propagation method

The beam propagation method (BPM) is based on one-way wave equations combined with a slowly varying envelope approximation [12]. This makes the simulation of the light propagation computationally efficient but limits the application of

this method to cases in which the variations in the refractive index are small. Because of the one-way assumption, BPM does not consider backward reflections, although some commercial tools have incorporated the capability to handle bidirectional propagation. The implication is that BPM is a good technique for designing and optimizing photonic devices that employ slowly bending waveguide structures, structures built in low contrast cross-sections, and structures where reflections do not play a major role. Based on the applications, it can be scalar, semivectorial or fully vectorial, and applicable to both 2D and 3D structures.

10.2.3.3 Finite difference time domain

While BPM is fast and simple, and EME includes natural support for reflections, the finite difference time domain (FDTD) method is the most versatile method in use to simulate photonic devices. FDTD solves Maxwell's equations directly on a discretized three-dimensional grid without any physical approximation [13]. Therefore, the FDTD method can simulate any effect described by Maxwell's equations, including scattering, transmission, reflection, absorption, nonlinearity, etc. However, it is a compute-resource heavy method and therefore the maximum problem size is limited by the extent of the computing power available. FDTD is a time-domain solution, and frequency response can be obtained by a single pulse or impulse simulation using Fourier transform techniques. Due to the requirements for compute power and memory usage, several methods to obtain accurate results faster have been developed, for example, clustering for utilizing computer farms and cloud installations, using radial or other symmetries and periodic boundary conditions.

In addition to these three main physical level photonic simulation algorithms, illustrated below in Fig. 10.5, one will find cosimulation techniques of combining ray tracing and rigorous solvers, or approaches that split a larger problem into subsequent sections that are solved with different rigorous methods.

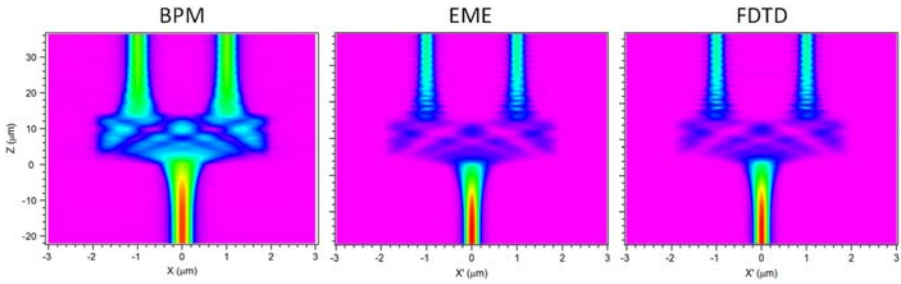

Figure 10.5 Top views of propagation simulations of a 1 × 2 multimode interferometer device, showing the lack of support for reflections with beam propagation method and very similar results for finite difference time domain and eigenmode expansion. Computational times go down from finite difference time domain = 1000, to eigenmode expansion = 500 to beam propagation method = 50 A.U.

10.2.4 Electro/optical codesign

To enable designers to design and optimize photonic devices that have electrical and optical behavioral characteristics, photonic device tools have been extended to incorporate multiphysics capabilities, supporting analysis of effects, such as refractive index changes as a function of electric field, temperature changes due to electrical current, or refractive index changes as a result of mechanical stresses. More complex are optoelectronic devices that rely on semiconducting behavior, like phase shifters, photodetectors, and light sources. In such devices, there is a close interaction of electrons and photons, coupled to dopant profiles and the resulting carrier distributions and electric fields.

To model such active devices, we can apply the same steps as for passive devices: starting from process simulation (now including dopant profiles and complex index profiles), the material properties can be fed into mode solvers and propagation tools to simulate the optical response. In the case of a photodetector, the optical field profile is fed back into an electrical device solver to derive the resulting photocurrent. A similar flow can be used to obtain the optical phase shift versus bias voltage for an electro-optic modulator. Often this electro-optical design is an iterative process between the various solvers. Streamlined workflows with a common cockpit and geometry database are commercially available.

A special case is the design of semiconductor amplifiers and lasers. While the basic physics of these devices are fairly well known, the actual behavior is heavily influenced by the precise composition of materials and boundary effects in the quantum wells, quantum dots or interfaces where unwanted recombination can occur [14]. Therefore, laser design is often done in specialized tools.

10.2.5 Feeding the circuit level design flow

The electromagnetic simulations needed to model photonic building blocks are computationally intensive, and scale with the size of the building block. When multiple building blocks are connected into a circuit, it is no longer practical to simulate the entire structure with these techniques. Instead, the circuit will be simulated by abstracting the building blocks into behavioral models that process the signals exchanged between building blocks. This is like the abstractions made in electronic circuit simulation.

The last decade has seen a strong push to bring integrated photonics design into the same realm as electronic integrated circuit (EIC) design and manufacture. The introduction of multiproject wafer or shuttle runs at institutes and later commercial facilities [15] has required the photonics community to start developing PDKs [5]. Optimized device designs, developed by device engineers, are collected, and bundled together in a PDK to bridge the information from the foundry to the design community, as shown in Fig. 10.6. To a circuit designer, the two most important aspects in this device library are the mask layout of the building block and its behavioral model. This behavioral model is a simplified model that is compatible with circuit level simulators, to allow photonic-integrated circuit (PIC) designers to analyze the performance of the overall circuit, as the physical or device level tools

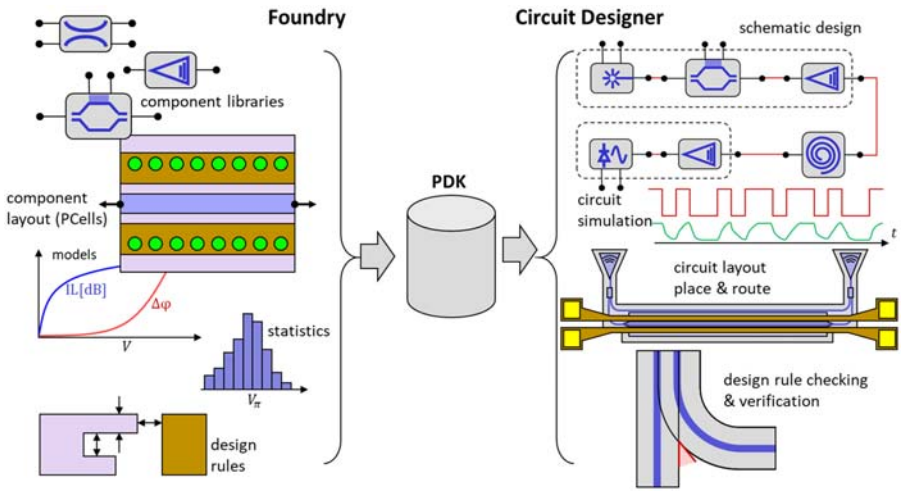

Figure 10.6 The process design kit is the bridge between the foundry or fab, and the circuit designer. It abstracts the actual processes in the fab by providing libraries of building blocks, models, and design rules.

are not able to calculate a meaningful result in an acceptable timespan. Photonic PDKs are discussed in more detail in Section 10.3.12.

One of the key outcomes of photonic device simulation is exactly the creation of such behavioral models. For passive devices, such models can be expressed as wavelength-dependent scatter parameters (S parameters) that can be directly extracted from simulations and confirmed or calibrated through optical transmission measurements of fabricated devices. For active devices, such as modulators, amplifiers, or detectors, the behavioral models are often captured in a set of equations that approximate the physics in the device, and for which the parameters are extracted from simulation and/or measurements. Alternatively, a mathematical model is fitted to the data using macromodeling or machine learning techniques.

As mentioned before, integrated photonic devices, like traditional semiconductor devices, are three-dimensional structures that are manufactured with a series of planar deposition and selective removal steps. The process and device simulators generate and use full 3D geometrical data, cross-sectional information, and top-view information. Today, there are only industry standard representations for the top-view representation, as this information is required for integration into a complete chip design and mask manufacturing. GDSII and OASIS are the two formats most used to carry information from device level design into full circuit design and from circuit design into final mask layout (discussed further in Section 10.3). For 3D representations and cross-sectional information, there are generic and tool specific formats that often require translation or conversion when designers want to utilize tools from different vendors. For photonics especially, the cross-sectional information is very valuable to support custom device design, a requirement that is, unlike in electronics, still a key activity for many organizations designing PICs.

In the early days of integrated photonics and still today when new fabrication processes or materials are introduced, behavioral models are derived from device-level simulations, until manufactured devices have been characterized. To automate the required steps, TCAD to SPICE tools exist for electronics and equivalent solutions are emerging for photonics [16]. These tools can run multiple device-level simulations varying material, geometric, physical parameters, and create a simplified model targeting photonic circuit level simulators.

10.3 Photonic-integrated circuit design

While photonic and electronic circuits operate on very different physical principles, the abstraction into circuits means that there are also many similarities: circuits consist of building blocks communicating through signal lines and can be organized into hierarchical subcircuits. It would therefore be a mistake to ignore the several decades of accumulated learning since the first EICs were commercially introduced in the early 1960s. In the early 1980s commercial CAD tools came on the scene for EIC design followed by the first commercial CAD tools for PIC design in the early 1990s. Just as the maturity of PIC fabrication processes lags that of electronics, the maturity of PIC circuit design techniques is two to three decades behind those for electronics, especially when it comes to the capability of guaranteeing first-time-right design [1]. But this gap is rapidly closing. The intersection of EIC and PIC CAD tools started to take place through a series of acquisitions and interoperability layers between CAD vendors in the late 2010s, establishing what is now commonly referred to as EPDA (electronic photonic design automation) [2]. Today, we see the first commercially available fully integrated electrical-optical CAD platforms that enable codesign of EIC and PIC designs.

The evolution of PIC design methodology follows closely that of EIC design over time. In the early years of both electronics and photonics, much focus was put on the physical aspects of how to best build and characterize the basic building blocks of the technology. In the last five years, the industry's adoption of silicon photonics, which makes use of existing CMOS manufacturing infrastructure, has greatly accelerated the pace at which integrated photonics has progressed, especially in terms of circuit scale and complexity. There is now a noticeable industry push toward circuit-level PIC design activity, much like what was experienced in the 1980s for EIC design. This is evidenced by the number of start-up companies creating PICs for applications including datacom, LIDAR, bio-sensing, AI, and quantum computing.

10.3.1 Formalisms

With the industry move to PIC design comes the maturing of how PICs are being designed. In the first two decades of EIC design, it was typical for one engineer to do all the steps of the design process, usually starting with the layout of the design.

Schematic formalisms were not necessary as the design complexity did not warrant it. The same has been true for integrated photonics. With the advent of application specific integrated circuit (ASIC) cell-based design methodologies in the 1980s came an explosion in the complexity of electronic designs. The industry shifted to schematic capture, circuit-level simulation, and schematic-driven layout (SDL) to handle the growing complexity. This same shift towards schematic-based design is now happening in the integrated photonics community.

Along with the shift in design methodology came a specialization of engineering tasks. Instead of one engineer doing all the tasks, the EIC was designed by a team of specialists. Front-end designers took on the tasks of schematic capture and simulation, while layout experts took on the task of assembling the integrated circuit and reticle. Formalism also took place in physical design verification and mask making, and there too engineers specialized in how to best do these tasks. The same separation of responsibilities is now happening for PIC designs, and where possible companies are sharing the PIC tasks amongst the specialists that formerly only worked on EICs.

It is yet to be seen if integrated photonics will make the same jump experienced in digital design with logic synthesis, optimization, and automated place and route replacing the SDL paradigm, or if photonics will follow the path of analog and RF design which still uses the SDL design approach today.

10.3.2 Process design kits and use of hierarchical design methodologies

A key item that accelerated the digital IC market forward in the early 1980s was the advent of library-based design. Previously, each transistor of an EIC was hand crafted. Library-based design offered a general set of basic characterized building blocks from which larger circuits could more easily be constructed. In the analog EIC markets these libraries became known as PDKs. The result of using library-based design is risk reduction owing to the use of prevalidated library components with known functionality and performance.

At the same time, commercial CAD tools started supporting hierarchical design methodologies that allowed designers to capture reusable logic as design blocks with the final leaf cells in the hierarchy being the library PDK components. Hierarchy enables multiple designers to work simultaneously on different parts of a design as well as at different levels of abstraction. This resulted in even more specialization, with separate teams being able to work simultaneously on PDK components, circuit-level functional blocks, full PIC die floorplans, and reticle floorplans. Hierarchy also allows for mixing higher system-level blocks together with the EICs to do system level verification, with each block being represented at different levels of abstraction.

Integrated photonics is now taking advantage of both the PDK library and hierarchy capabilities, and in fact system design now employs electronic–photonic (E/O) codesign in the same hierarchical CAD flow. A deeper dive into what makes up an E/O PDK is described in more detail in Section 10.3.13.

10.3.3 Design flow: from idea to manufacturing

While design of the EICs and PICs are different due to the nature of their respective domains, CAD flows have been developed to create solutions covering the various steps of the design flow. There are, however, several differences between electronics and photonics that necessitated changes to the standard EIC design tools and flow. The most notable differences are:

- Any-angle curvilinear layout shapes in photonics versus rectilinear Manhattan shapes in electronics.
- Only one or two interconnect layers in photonics versus multiple interconnect layers in electronics.
- Photonic signals can be allowed to cross through each other, using a crossing, versus this causing a short in electronics.
- Photonic waveguides propagate signals as waves which incur a time delay and phase delay during the propagation. In electronics, all points on a net are considered at the same potential at any given time.
- Photonic waveguides carry different signals in both directions and in different propagation modes.
- Photonic waveguides can carry multiple signals on different carrier wavelengths at the same time.
- Photonic waveguides only connect two ports on the circuit cells, while electronic wires allow multiple fanouts.
- Different waveguide interconnects can be built on the same layer versus single type in electronics
- Design for test (DFT) and DFM for integrated photonic ICs are not yet as mature as for EIC design

The typical PIC design to manufacturing flow (Fig. 10.7) is made up of the following steps that progress more or less in linear order [1,2]. With that said, designers will iterate back and forth between these steps as the design is defined and refined.

1. Schematic design capture;
2. Pre-layout simulation;
3. lSDL;
4. Back-annotation and post-layout simulation;
5. Physical verification (manufacturing design rule checks (DRC) and layout vs. schematic (LVS));
6. Post-layout data prep (Boolean operations, biasing, fill and optical proximity correction);
7. Design submission to the foundry (GDSII or OASIS tape out); and
8. Manufacturing, packaging, and test.

As discussed in Section 10.2.5, this design workflow is fed by a PDK provided by the foundry that will manufacture the PIC. Due to the current immaturity of integrated photonics or given specific end application requirements, the designer may need to create additional custom components that are compatible with the process technology used by the foundry PDK. In this case both custom components and PDK components are instantiated in the design, connected with waveguides,

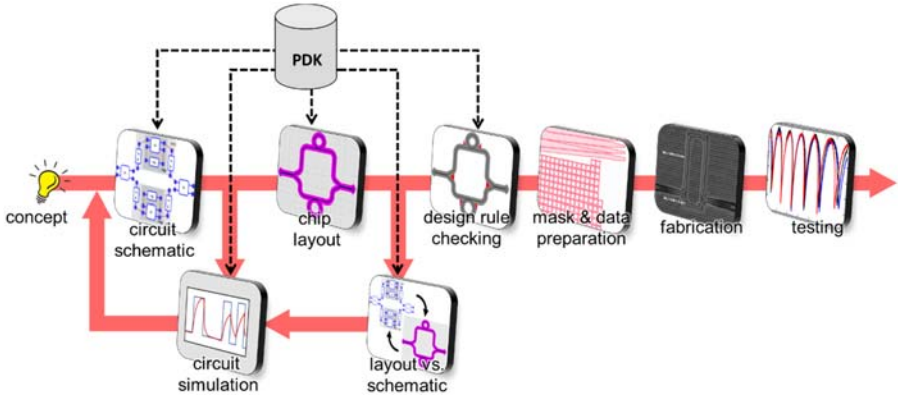

Figure 10.7 Schematic driven layout flow, using a process design kit with symbols, simulation models, (parametric) layout and physical verification decks for design rule checking and layout versus schematic validation.

simulated, and laid out. This results in a risk reduction due to verified performance of the PDK library and offers a platform that the designer can use to accelerate PIC design to create a highly optimized end design/prototype.

As integrated photonic design matures, we are also likely to see the emergence of additional steps that focus on DFM and DFT.

10.3.4 Schematic capture

As with EIC design in the 1980s, there are today PDKs for photonic circuit design, but there is not yet an automated way to convert abstract system-level behavioral representations of photonic designs to functional circuits based on the PDK components. Instead, designers manually map the photonic design from abstract system-level representations to an equivalent circuit assembled from the basic photonic PDK building blocks using a schematic capture CAD tool. This works well, as PICs are typically integrated with custom analog electronic circuitry and these types of electronic circuits also use a schematic driven design methodology. The CAD tools are flexible enough to allow for scenarios where the electronics and photonics are on a different die or monolithically integrated on the same die.

The CAD flows should differentiate between electronic and photonic connections. Design cells can be purely electronic, purely photonic or a combination of the two. Schematic checks are employed to ensure designers do not accidentally connect photonic pins to electronic pins, and fanout checks ensure that no more than two photonic pins connect to a single waveguide. As in custom analog design, there is often a need for large numbers of variants of a given photonic cell and as such the cells are typically created as parameterized building blocks, commonly referred to as PCells. Electronic and photonic connectivity is defined by connecting pins that represent the logical connection points on the cells, using schematic wires

to indicate logical connections between these pins. Cells are instantiated in the schematic from a PDK library menu, and the cell parameters can be set on each individual instance.

Because photonic waveguide connections can use waveguides of different types (e.g., with different cross-sections, such as a partially etched rib or a fully etched strip), and because the physical characteristics of the waveguides have a big impact on the photonic circuit performance, the photonic waveguide connections are also implemented as programmable photonic building blocks (PCells) with parameter settings for the waveguide width, type, and other geometrical properties, such as incident angle, rotation, length, and radius of curvature for arcs and bends. During the SDL step, these parameters are passed on to the layout tool to ensure that the designer's intent is preserved in the layout.

In the simplest case, the designer instantiates two photonic cells and one connecting photonic waveguide. Unlike electronic connections, which are often considered as logical, instantaneous wires, the connecting waveguide will also be represented as a building block, as it will induce a propagation delay and loss to the photonic signal and should therefore also be modeled as a block with these optical properties. Thus, the output pin of the first photonic cell is connected by a schematic wire to the input pin of the photonic waveguide, and the output pin of that waveguide is then connected by another schematic wire to the input pin of a second photonic cell. In this approach, the schematic wire represents a perfect, zero-length optical connection between the waveguide ports of the building blocks and the waveguide. An arbitrarily complex circuit can be captured using this method when employed in a hierarchical approach, where small circuits become the building blocks of more advanced circuits.

10.3.5 Circuit simulation

Before proceeding to the layout phase, designers may simulate the PIC at the circuit level using the default or ideal component parameter values or the values as defined by the designer. Even though light waves are electromagnetic waves and could in principle be mapped onto voltages and currents, such mapping is not convenient for a designer [17]. Instead, as discussed in Section 10.2.5, the behavior of the building blocks is modeled as a scattering of incoming waves into the outgoing waveguide ports [18,19]. Each waveguide can also carry multiple spatial modes, polarizations, and carrier wavelengths [4,7].

Section 10.3.11 provides more detail about the difference between photonic and electronic circuit simulations; here, it suffices to say that it is very difficult to accurately capture photonic behavior in an electric simulator, and it is preferred to use a specialized photonic circuit simulator for that purpose. When combining photonic and electronic elements in a single circuit, a cosimulation solution is needed. Such a mixed-signal simulation requires a linking of photonic and electronic circuit simulation engines, exchanging information through the opto-electronic blocks in the circuit [20,21].

Upon invoking a circuit simulation interface, the designers can set up virtual test benches detailing the various types of input stimulus and analysis routines to be performed, and they can set probes at any point in the circuit schematic to monitor the electrical or optical signals at those points. Upon invocation of the simulator, the CAD system automatically creates the appropriate connectivity netlists of the circuit for the electrical and optical simulators being used, including the parameters as set by the designer, and then runs the simulation. Simulation results are viewed in a waveform viewer where additional analysis can be done on the probed signal waveforms.

It should be noted that in a hierarchical design flow, designers can simulate individual levels of hierarchy separately or in context of the larger design. Not all parts of the design need be completed before the simulation phase can begin.

10.3.6 Schematic driven layout

Once the functional schematic is ready, the physical design can proceed using an SDL approach. SDL represents a connectivity-driven approach where, for each instance in the schematic, a PDK layout PCell is selected and placed on the layout canvas. As cells are placed, flight lines indicate the connections that have yet to be made. Connections are done by abutting waveguide ports, or through automation aides, such as interactive connectors and auto-routers. The designer iterates on placements and connections until all components from the schematic are placed and connected.

Note that defining optical connections can change the functionality of the circuit. Optical waveguides, either manually or automatically routed, introduce a phase delay and time delay for the optical waves. Depending on the desired functionality, these delays can affect the performance of the circuit, especially when they are used in wavelength filters. Therefore, it is important that the design intent (e.g., whether a link is phase sensitive or not) is properly defined in the schematic and correctly carried over to the layout.

If the designer is working on a hierarchical block, hierarchical pins are placed as connection points for when the cell is instantiated later in another layout. Since the schematic designer does not see the physical context into which components are being placed, it is likely that the layout engineer will need to update some of the component parameters to make the design fit in the area given. This is especially true for the waveguide elements, as these are continually manipulated until the layout is complete. CAD flows are flexible and allow for different policies to be enforced during the SDL process. Some parameters may not be editable in the layout editor and only the schematic designers can change them, while other parameters can be changed directly by the layout engineer. Changes made by the layout engineer can be back-annotated to the schematic for use in postlayout simulations.

Because the photonic waveguide layout can dramatically affect the schematic, some design tools merge these two steps, allowing the designer to define the schematic connectivity directly in the layout [22]. This can reduce the back-and-forth iteration between the schematic and layout design, but it also introduces the drawback that a design might not be easily ported between different fabrication processes or foundries.

10.3.7 Back-annotation and postlayout simulation

After the circuit has been laid out, it is important to verify whether it matches the original schematic and performs as intended. For that, the layout should be translated back into a circuit model that can be simulated, but now with the actual parameters of the blocks used in the layout. The common way to do this is to use the information in the layout to back-annotate the circuit as it has been defined in the schematic.

As there is a fundamental difference in how electrical and photonic connections are treated in CAD flows today, there are obviously similar differences in how the CAD flows back-annotate these changes for postlayout simulation. For most electrical connections, the frequencies of the signal are slow enough to allow the tools to make simplifying assumptions about the parasitic effects of the layout connection (the metal wires) on the electrical signal. A tool known as a parasitic extractor is run on the electrical connections of the layout in a 2D or 2.5D flow to create an equivalent RC (resistance−capacitance) network of the interconnects that is back-annotated to the schematic database for use during electrical simulations. At higher speeds, the frequency dependent nature of this RC network must also be calculated. In this case, a 2D or 2.5D analysis is no longer adequate and a full 3D solver must be employed. Additionally, electrical parasitic extractors also look at the interaction between neighboring interconnects to determine if there is crosstalk between the electrical signals.

Currently, in the photonic realm, an equivalent of parasitic extraction does not yet exist. Instead the waveguides are treated as parameterized cells that are characterized by various parameters, such as their widths, waveguide type (rib, strip, ...), and radius of curvature for bends and curves. The final waveguide path between components in the layout is back-annotated as a network of waveguides with appropriate parameters. The simulation models then pick up the back-annotated parameters and the full netlist with all its constituent components to perform the postlayout simulation. A gap in this methodology is that there is not yet an analysis of the interaction between waveguides that are (too) close to each other. As PIC design flows mature, it is envisioned that the equivalent of electrical parasitic extraction will occur for photonic waveguides as well.

It should be noted that a true electro-optical codesign CAD flow should enable both electrical and photonic back-annotations to occur, for accurate postlayout cosimulation of both the electronic and photonic parts of the circuit. In design flows where the photonic schematic is defined directly in a layout environment, the layout-dependent properties are immediately incorporated into the circuit simulation.

10.3.8 Physical verification (design rule and layout vs. schematic checking)

Once the layout has been generated, it must be checked for possible design errors that could result in a defective chip. Since both EICs and PICs use similar wafer processing techniques it makes sense to use the same EIC methodologies and CAD tools for PIC

DRC. Conventional DRC tools are being used for PIC manufacturing layout checks, and in fact, they are required by most foundries before manufacturing of a submitted design is started. Foundries are responsible for writing the DRC rule decks and special care need to be taken to develop robust checks for curvilinear PIC layouts [23,24]. These DRC decks become an essential part of the foundry photonic PDK. Typical DRC rules check for minimum feature widths and spaces, sharp angles, minimum overlap, or separation between patterns on different process layers, and pattern densities.

In addition to layout DRC, some foundries and CAD suppliers are now offering photonic layout versus schematic checks. In electronic design, the LVS step checks that the layout is connected as specified in the schematic. LVS also extracts various parameters of the design, such as L/W for transistors, and then compares those extracted parameters against parameters as captured in the schematic. Photonic LVS follows in a similar vein, first by doing photonic light path connectivity checks and then by extraction of parameters from the layout for comparison against those found in the schematic [25]. As PIC designs become more complex, a combined E/O LVS strategy is required to ensure both the electrical and photonic paths are configured and connected correctly.

10.3.9 Design for manufacturability, design for test, and post-layout data prep

When a device is fabricated, every process step will introduce some variation in layer thickness, line width or material composition. When designing photonic circuits, it is important to consider that these circuits should be manufacturable, that is, that the circuits will work even in the presence of these variations. Section 10.5 covers the need for more advanced DFM and DFT methodologies. This includes the use of stochastic data as measured from the fabrication of many wafers. This data can be used in the aforementioned design flow using Monte-Carlo type simulations [6,7]. This typically is done during or after the postlayout simulation phase. However, quite often Monte-Carlo simulations are also employed after the design has been fabricated, to help designers center the design in the manufacturing process window during subsequent revisions of the design.

It should also be noted that normal EIC post-layout data prep functions are also required for PIC design. PIC design is especially complex in that many applications require heterointegration of multiple photonic dies, sometimes bonded to each other. This drives the requirement for the CAD tools to create layouts of different mask sets to be used in multiple foundries. Each foundry has its own biasing and sizing operations based on their process, and care must be taken that the right postlayout operations are done on each layer's data before taping out to the respective foundry.

10.3.10 Manufacturing, packaging, and test

With few exceptions, mask and wafer manufacturing for PICs follows closely that of EICs. Packaging and test however take on a variety of new complexities.

These are covered in Chapter 12. The main point to be made here is that since manufacturing variances play such a large role in the functionality and performance of PICs, designers would be wise to pay close attention to closing the loop between actual measured PIC performance versus performance as predicted by simulations.

Although basic photonic elements can be sufficient for a custom design tape out, a fully functional design requires expertise in process, design, test, and layout. Today, a photonic design typically takes multiple fabrication runs (iterations) to meet the full set of specifications, which increases the cost of development and slows down the development of a product. The added cost and development time can lead to a loss of competitive advantages that come with the custom design. Instead, a PDK component library aided design is preferred to accurately predict PIC behavior and iterate at a rapid pace, which reduces the risk and time to market. In addition, the overall cost of licensing a component library can be less than the custom development of those components. Most successful companies designing PICs today plan for multiple iterations through the foundry to tune their PICs, much like analog design of the 1980s. It should also be noted that attention should be paid up front to how the PIC will be tested in a volume manufacturing environment. Again, this is much like testing of analog electronics.

10.3.11 Notes on electrical versus optical simulation and modeling

As mentioned earlier, while the design flows for electrical and optical ICs are similar, there are some areas of very distinct differences that must be clearly understood. One of these is in the modeling and simulation of photonic devices and circuits.

We can distinguish two distinct types of photonic circuit simulations: frequency domain and time domain. Frequency-domain simulations are used to characterize the wavelength-dependent transmission of a circuit (e.g., a wavelength filter) for a given static configuration. This can give useful information about wavelength dispersion or crosstalk between different wavelength channels in a WDM link. The frequency or wavelength response is either calculated iteratively or through a numerical matrix-inversion procedure [18]. The resulting wavelength-dependent scatter matrix can be persisted through different file formats (e.g., Touchstone).

To obtain the time-dependent behavior, a time-stepped simulation is used where the optical signals are passed between the components. These optical signals differ significantly from electrical signals, as summarized in [4]. Electrical signals are real-valued, base-band signals, and their Fourier transform is Hermitian, described by the positive frequency part of the spectrum. In contrast, optical signals are narrow band with a complex-valued (real and imaginary) envelope that describes the signal amplitude and phase; thus optical signals are represented as a complex envelope with a center frequency. Optical signals can be single-wavelength or multiple-wavelength, and they have polarization states and transverse mode profiles. There simply is no equivalent of these optical attributes in the electrical domain. Note that

we mean with "narrow-band" that the bandwidth of the optical signal is small compared to its optical carrier frequency (180−240 THz), but this "narrow" bandwidth of each wavelength channel is still large for electronic circuits ($\sim$100 GHz).

The flow of optical signals through a circuit is described by the scattering of the carrier waves in each building block. This scattering can be described by a static matrix, analytical equations or a macromodel fitted to electromagnetic simulations or measurements. The scattering model can also be dependent on the voltages or current in the electrical pins, or even capture nonlinear optical effects. It is important to point out that optical signals are inherently bidirectional and the scattering in components can give rise to reflections. Whether intentional or not, such reflections significantly affect circuit performance, as they can induce in-circuit resonances, inter- and intrachannel crosstalk, polarization crosstalk, phase- and intensity-noise, chromatic- and polarization-dispersion and nonlinearities. All of these are very difficult to model in an electrical SPICE simulator and as such, this method of modeling and simulation should be avoided in favor of using dedicated photonic simulators that are natively built to model all these phenomena.

Still, there are examples where electrical simulators are used to model photonic circuits, using Verilog-A to define photonic models [26,27]. However, this is not a universally applicable method and there is a time and place for using Verilog-A or electrically equivalent models for optical subcircuits. Good examples of this are when the designer is working to optimize the electrical circuitry that will drive a modulator or the receiver circuitry that will convert current from a photo diode back to the electrical domain. However, when modeling the optical circuits or cells to be used in the PIC, true photonic models should be employed. An example is a simple set of back-to-back gratings with a waveguide connection between them [4]; when broadband light is applied as a source, multiple orders of coherent reflections result in a high-Q narrow passband filter. If the models do not support reflections, the narrow band of the filter is totally missed, and the design would erroneously appear to be a band-reject filter, as can be seen in Fig. 10.8.

Dedicated PIC simulators do support different combinations of photonic signals, but at this time there is little or no standardization. Circuit models need to be constructed or optimized for each circuit simulator. Simple models (e.g., based on static scatter matrix) can be recompiled for different simulators, but models based on equations or code need to be redeveloped, especially if they make use of unique features of one particular circuit simulation engine. Some standardization of behavioral photonic model definitions would be a strong enabler for the PIC designer and industry.

As already mentioned in Section 10.3.5, photonic−electronic cosimulation requires a mixed-signal approach with a photonic and electronic circuit simulation engine [20,21]. Fortunately, the slowly varying envelope representation of optical signals eliminates the need for Terahertz rate sampling, enabling cosimulation of both electrical and optical signals with a common time unit resolution at modulation speeds.

To summarize, the PIC designer should use the right tool for the job at hand. In electronics, different model abstractions are used for different purposes (transistor

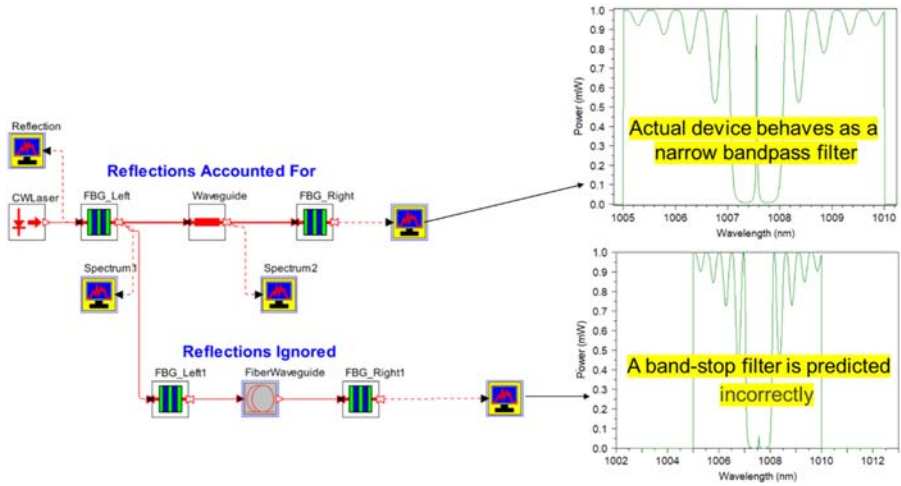

Figure 10.8 Simulation results of a simple circuit with back-to-back gratings connected with a waveguide or a fiber model with and without considering reflections of the propagating optical signal.

level, gate level, behavioral level, etc.). The same rule applies for photonics. CAD flows are now commercially available at these different levels of modeling abstraction, along with the capability to do multidomain cosimulation with the different simulation engines and their respective models.

10.3.12 Process design kits

A key enabler for complex photonic circuit design is the PDK. As depicted in Fig. 10.6, it is essentially the bridge between the fab/foundry and the circuit designer, as it provides the necessary abstractions that enable the circuit designer to capture, simulate and layout a design, and predict how it will function [5].

High quality PDKs enable the creation and rapid development of world-class silicon photonics. PDKs encapsulate information gained through multiple iterations of design, fabrication, and characterization cycles. These typically contain design guides, technology files, DRC decks, custom design workflows incorporating EPDA methodology support from the foundry and multiple component libraries (documentation, abstracts, schematics, compact models, statistical models, and system level simulation deliverables in the major CAD flows) from both the foundry and third-party suppliers. PDKs should also include additional documentation in the form of a Packaging Design Guide and Packaging CAD views to support test, assembly, and packaging.

A unique feature of photonic PDKs is that the components of the PDKs must be tuned to target end-applications that imply a range of wavelengths. A photonic component design for telecom wavelengths in one band may not function well or at all in another band. This means that the foundries must pick and choose which

end-application spaces they will support with their PDKs, and then tune the PDK components accordingly. Some foundries are already working to support multiple PDKs so that they can engage in the very different end application spaces, such as telecom/datacom, AI/ML, LIDAR, and bio-sensing, to name a few. As part of these PDKs, the foundry team may also bring in photonic PDK components from third-party sources that have been targeted to a given end-application space and the foundry's process.

Typical photonic PDKs are made up of the follow elements:

- Documentation:
  - o Design guides;
  - o Component library documentation;
  - o Alternate component library documentation;
  - o PDK overview—overview and release notes.
- CAD views (schematic symbols, simulation models, layout cells, connectivity rules, parameter definitions and defaults and call backs);
- Technology files (technology stack, material stack data with photonic parameters, such as effective index, target wavelength, and basic photonic design rules, layer maps used to create GDS or OASIS formats, display files with standard color and patterns per layer);
- Stochastic model data supporting statistical simulations;
- Physical verification decks for DRC and LVS tools;
- Passive models (S parameters for passive components provided for users without access to CAD views for supported software tools); and
- PDK component library GDS (component library abstracts in component and GDS formats provided for users without access to layout through CAD views for supported software tools).

Note that today, not all foundries provide the same level of detail in their PDK, and not all design tools can handle all this information. As already mentioned, a significant obstacle is the lack of standard circuit models for photonic components, which implies that PDK builders need to either construct their building block models in different tools, or limit the PDK to a subset of the available design tools.

10.3.13 Process design kits for monolithic electro-optical processes

Most foundries support separate processes for EIC and PIC designs. However at least one major foundry is supporting a monolithic process in which active electronic and photonic devices can be placed on the same die [28]. In this case, the PDK becomes a superset of the traditional analog/mixed signal PDK and the photonic PDK. Multiple libraries of building blocks are included for the electrical devices and the photonic devices. Each library contains the various CAD views (symbol, layout, simulation model, LVS, etc.) needed for the electrical and photonic parts of the flow. Extra data is also included for the electrical parasitic extractor(s) supported in the flow, and extra DRC and LVS decks and runsets are provided to complete the physical verification flow. The PDK may also have additional interface elements and models to enable codesign in the CAD flow. The major

difference between this type of PDK and a photonics only PDK is the amount of flow testing that must take place to ensure that the constituent pieces all follow the same CAD methodologies that enable the monolithic process to work.

One last note is that PDKs are now also starting to add components for E/O codesign that consider package configurations. This would include standard package templates designating both electrical and optical interface locations, pad configurations and the like.

10.4 System design

The design of the photonic chip cannot be considered as isolated from the design of the entire system or subsystems in which the chips will be used. The chip design is therefore just a part of the system and subsystem design. An optical data link, from the point of view of the system, transmits and receives a digital data stream over one or more high-speed digital signal lanes. In addition, the optical link could have a (low speed) digital interface to configure and monitor its properties.

The photonic chip, as part of such a fiber optic link, is interfaced with multiple layers of electronics. The digital signals coming into the transmitter need to be switched, serialized, and converted to an analog modulation signal for the electro-optic modulators. On the receiver side, the optical signals in the photodetector need to be digitized, deserialized and switched. Apart from these high-speed electronics, the photonic chip needs to be controlled, for example, for maintaining alignment with the transmission wavelength. The design of the photonic chip and these driver electronics cannot be fully decoupled.

The photonics and electronics need to be physically integrated. Either this is done on the same chip, using monolithic photonic−electronic integration, or the photonics and electronics are designed on separate chips. In both cases, there is a need for codesign with the package to handle the optical and electrical signals off-chip. When photonics and electronics are implemented in separate chips, there is also a challenging high-speed RF design at the interface between the two.

As touched upon in Chapter 12 the package design is a true multiphysics problem, combining optical design (laser integration, fiber attachment), electrical design (PCB, interposers and connectors), thermal design (self-heating, thermal crosstalk, material aging, thermal control), mechanical design (stress, thermal expansion, shocks), RF design (strip lines, connectors), and a study of the electromagnetic interference. The design also needs to fit the constraints of the physical form factor of the pluggable module, on-board module, or near or copackaged optics used in the overall datacom system.

There is no standardized system design flow that incorporates all these considerations. For each problem there are suitable CAD tools that can help an expert to evaluate the electrical, thermal, optical, and mechanical performance, but the design decisions are often based on “smart choices” by experienced system designers. This practice is unlikely to change soon, especially as the technological evolutions are

moving at such a rapid pace that there is little opportunity to standardize these aspects of the design flow.

10.5 Design for manufacturability

With its high refractive index contrast, silicon photonics enables dense packing of submicrometer optical waveguides and photonic components. As a result, PICs are reaching a level of complexity such that several hundred, or even thousands, of individual components in the PIC must function to meet demanding specifications, in order for the entire PIC to have satisfactory performance and yield. This demands that manufacturing variations affecting photonic components be considered and addressed as much as possible during both the photonic component/device design stage, and the circuit design stage. That is to say, DFM becomes an essential component of the photonic device and circuit design methodology. This is especially critical in silicon photonics because: the same high refractive index contrast that enables large-scale integration also makes the waveguides and building blocks sensitive to nanometer-scale deviations in the fabricated geometry.

In this section, we describe the key elements of emerging photonic DFM concepts and methods. These include understanding of the impacts of manufacturing variations on geometry and material properties, the scale and scope of such variations, and the random versus systematic nature of different variations. We then describe ways to construct and use such variational models, ranging from simple corner models to parametric or variation-aware compact models. Taking advantage of these models, approaches for yield analysis and optimization are then discussed. Finally, we emphasize the need in the future for not only DFM, but also photonic DFT.

10.5.1 Manufacturing variations

Manufacturing variations comprise any deviation in the intended geometric structure or its material properties, arising as an inherent limitation of the fabrication process. Because integrated photonic devices and components generally consist of multiple thin films (spanning dielectrics, semiconductors, and metals) that are deposited, patterned, etched, doped, oxidized, or otherwise formed and changed across hundreds of process steps, a wide range of different kinds of manufacturing variations can be of concern. One major category is geometric variations, including deviations from intended film thickness (and thus height of elements of photonic structures), as well as deviations in patterned widths or other dimensions, including sidewall angles. Alignment variations due to overlay or offsets of one layer to another can also be important. For high-contrast waveguides as in silicon photonics, the fabrication variations that get the most attention are those that affect the interfaces between silicon and the surrounding oxide. Nanometer-scale variations in the width of a waveguide or the thickness of the silicon film can result in a

nonnegligible change in behavior of a waveguide circuit. In addition to geometric variations, material properties often need to be well controlled, from doping in semiconductor structures to composition in dielectric or waveguide structures, as these can affect optical properties, such as the index of refraction and loss parameters.

Just as in electronic ICs, manufacturing variations can have different statistical properties with random or systematic manifestations, often including spatial components [29–31]. Some variations are highly "random" and local in nature. For example, a spot or particle defect may appear with some probability at any location on a layout and with some probabilistic size distribution (e.g., a metal particle being embedded within or near to a dielectric waveguide). An important concern for low-loss waveguide and related photonic components is line edge roughness (LER), typically characterized as a short-range spatially correlated variation along the sidewalls of patterned structures. Other variations may have even longer-range spatially correlated deviations. For example, the thickness of the silicon layer in silicon-on-insulator (SOI) technologies, often intended to be near 220 nm, can deviate by several nanometers but does so slowly across hundreds of microns (within a chip) or millimeters (across the wafer). Finally, some variations are known to have systematic layout dependent components, often associated with layout pattern density [32]. For example, plasma etch rates and thus etched feature thicknesses, widths, or shapes can depend on the pattern density (area fraction of exposed film) being etched. Similarly, the chemical mechanical polishing (CMP) rate and thus structure thicknesses can depend strongly on the layout pattern density (raised area fraction) [33]. While dummy fill design rules are commonly included in foundry PDKs to bound or limit such systematic variations, an appreciation of these systematic variations remain important for proper function of some photonic circuits.

10.5.2 Manufacturing variation-aware photonic models

A variety of methods (Fig. 10.9) are typically used to understand and model the impact of the variations described above. In general, these models build on and go beyond the "nominal" models of photonic devices and circuits discussed previously

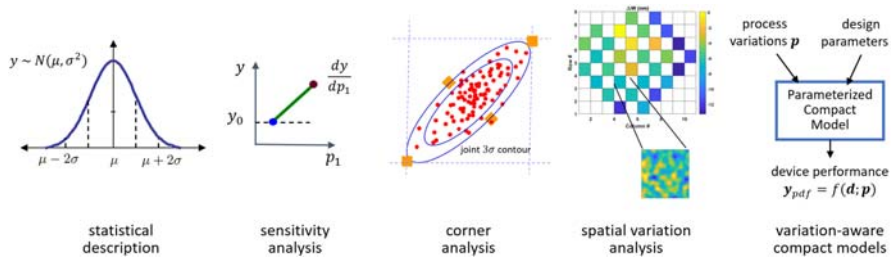

Figure 10.9 Range of variation-aware models, from simple (aggregate Gaussian statistical) to more sophisticated (corner models; sensitivity models; spatial correlation models; parameterized compact models).

in this chapter, and seek to provide information about the range and nature of device/circuit performance that might be encountered due to these variations.

Perhaps the simplest, but still highly valuable, variation model is the expression of a key geometric or device parameter as a probability distribution with associated statistics, for example, as a gaussian distribution with standard deviation (σ) in addition to the nominal or mean value (μ). Such aggregate variation measures (e.g., for layer thicknesses in the layer stack, or for etched feature widths), while limited, must be demanded to be provided by foundry and fabs by designers; without such information, high-yield design of increasingly complex PICs becomes nearly impossible.

A relatively straight-forward and often useful form of variation-aware models that complement basic mean/sigma models, are "corner models" that seek to establish the outer limits of variation source or device performance parameters. Each corner model, conventionally labeled with "high" (H) or "low" (L) deviations from nominal (N) for the corresponding parameter of interest, are intended to bound the expected range of deviations (e.g., representing three-sigma limits). An advantage of a corner model is that it represents a device instance having those deviations using a "standard" photonic compact model description (e.g., port-based S parameters) for that instance. In electronic PDKs where transistor speed is a good overall indicator for NMOS and PMOS devices, a small number of fast/slow corners are typically included in PDKs. Unfortunately, there is not yet a generally agreed upon set of key characteristics for the range of photonic components, and thus photonic PDKs may contain or provide different corner models, if they provide corner models at all. Also, for some model parameters there is no relevant classification as "good" or "bad," but instead it is important that the parameters of the building blocks are closely matched within the same (sub)circuit.

Another relatively straight-forward variation impact method is sensitivity analysis. In this case, typically first order (linearized) impact of some set of variation sources can be analyzed, typically by way of simulation, or sometimes by way of differential test structure design and measurement and provided to the designer. For example, the impact of through-port transmission T due to small deviations in waveguide thickness x might be expressed as $\mathrm{d}T/\mathrm{d}x$. Combining this with expectations or models of waveguide linewidth and thickness variations, distributions over the wafer can be used in subsequent analysis to estimate expected variations in transmission.

Generalizing the idea of sensitivity analysis are parameterized variation-aware compact models [34]. Such models can express, for example, the S parameters for a photonic component as a function of not only wavelength, but also as a function of key geometric (or material property) parameters, or their deviation from the nominal value. Thus, in yield analysis and optimization as discussed later, statistical simulations can be efficiently considered to understand the photonic device-level and circuit-level impact of many combinations of manufacturing variations. Variation-aware compact models can benefit from sophisticated statistical and reduced-order modeling methods, coupled with efficient sampling using more costly fundamental simulations (e.g., FDTD), capturing and enabling expression of

both random variation effects (like LER), and geometry-dependent variations (layer thicknesses, feature widths, gaps, etc.) [35].

Finally, spatial or layout-dependent variation models require additional modeling methods. Extending beyond just an aggregate mean/variance description of some parameter, a spatial model provides additional statistical components, most commonly the spatial correlation length over which the parameter varies. For parameters that vary slowly over the entire wafer, the spatial correlation length may be tens of millimeters; for other more local parameters (like LER), the correlation length may be on the order of micrometers. Other parameters, such as those associated with pattern density dependencies, may include a systematic component in either simplified form (e.g., sensitivity-analysis like) or chip-scale variation analysis form (e.g., a simulation of thickness variation across the chip, as a function of the specific layout of the chip) [6,7].

10.5.3 Yield analysis and optimization

Yield analysis involves propagating a description or model of an underlying variation (e.g., a model of film thickness variation) through to the impact on a performance characteristic of the device or circuit, to estimate what fraction of fabricated devices or circuits will function satisfactorily.

Corner-based analysis at the device level is one example. In this case, the device design might be simulated at the nominal and each of the corners; if the device performance is within specification at all of these points, and these points span 95% (say) of the true range of those variations, then the hope is that the yield will be at least 95%. While this is a good starting analysis, designers should be cautious about each step in this chain of logic. Uncorrelated corner models may be overly optimistic or overly pessimistic when the variations in fact have correlations. For example, if parameter p_1 and p_2 in Fig. 10.10 are correlated, but univariate limits are used, the design may appear to fail at the LH corner simulation, when in the actual fabrication process the LH corner will essentially never arise.

To handle correlated process variations, Monte-Carlo or other related statistical simulations are often used. In this case, samples of parameters (e.g., $p_1$and p_2) are generated based on the correlated (and/or spatial) statistics of those parameters, and photonic device or circuit simulations are performed to estimate the statistics (or at the very least, the yield) of the resulting performance parameters. Because substantial numbers of simulations must be performed, such statistical simulations demand and benefit from variation-aware models that can produce compact models for circuit-level photonic simulation.

Spatially correlated process variations can be particularly important for photonic circuits. In some cases, if all of the components in a circuit experience the same deviation (e.g., are highly correlated due to die-to-die variations), then the photonic performance may be minimally perturbed (e.g., waveguide modes remain well matched). In other cases, such as in coupled resonator optical waveguides that can extend spatially over long ranges, within-die spatial deviations in geometric parameters can destroy resonance and circuit operation. Analysis of the effects of

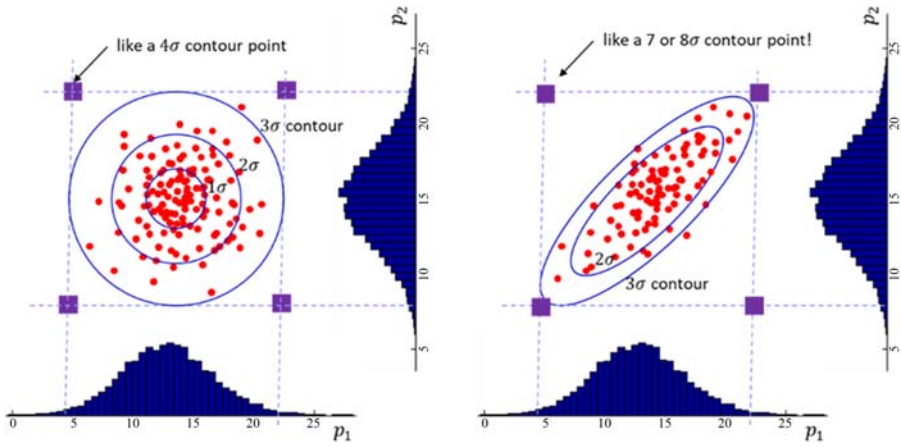

Figure 10.10 Left: Uncorrelated (univariate) corners analysis; corners may represent, for example, 4-sigma likelihood. Right: If the process variations are corrected as indicated by scatterplot, then some corners will be overly pessimistic, for example, corresponding to 7- or 8-sigma likelihood.

spatially correlated and layout-dependent process variations motivates the development of variation-aware compact models, coupled with variation extraction and decomposition methods [36]. Additional methods for analysis of the impact of particle defects on photonic devices are just emerging [37]; as the density and complexity of photonic circuits increase, such approaches may become important in the future.

Finally, DFM also includes methods that seek to optimize the design, to achieve both high yield and high performance. It should be noted that there is generally a tradeoff between yield and performance. Purely performance-driven design that proceeds by finding a design that maximizes an output performance and sits right up against a process limit (as pictured in Fig. 10.11), will almost always be a low yield design because many of the outputs will deviate from that process limit. Adding margins is a conventional approach to push the nominal design away from these precarious process limits; such margins can be approximated based on simple statistical models coupled with sensitivity analysis and corner models. For more complete descriptions of manufacturing variations and more complex nonlinear interactions with the photonic device or circuit, design centering approaches seek to find a nominal design that will still yield a large fraction of fabricated circuits in the face of manufacturing variations, with Monte-Carlo or statistical simulations [38].

A related problem in photonic circuits is robustness to operating condition variations, in addition to manufacturing variations. In electronic ICs, margins and statistical design are often applied broadly to guard against "process, voltage and thermal" (or PVT) variations. Similarly, in photonic devices and circuits, small variations in thermal or other operating conditions can sometimes be accommodated

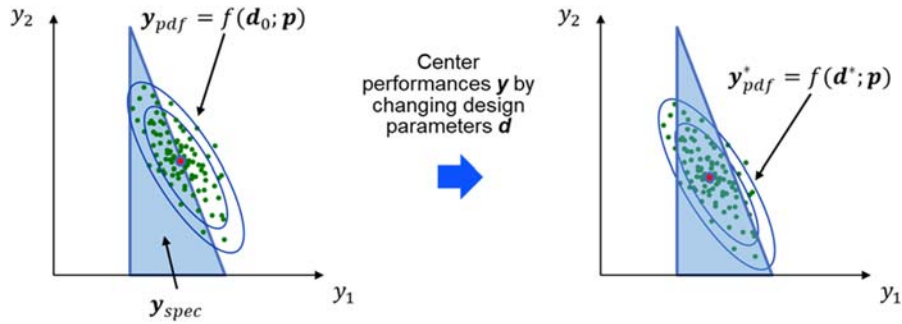

Figure 10.11 Nominal design (left) with low yield versus design centering (right) with high yield.

with design margins or design centering. Often, however, thermal variations in particular can have large effects on photonic devices and circuits, requiring active control circuitry and compensation to maintain resonance and other critical operating points [310].

10.5.4 Photonics design rule checking, FILL, and optical proximity correction

The designer of photonic devices and circuits will also benefit from stages in the design flow that are inherited, sometimes with some modification, from electronics CAD and DFM. An important stage is DRC described in Section 10.3.8 that involves checking of the layout for violations of design rules generally provided by the foundry.

One technique to improve the fidelity of the fabrication with respect to the original design patterns is the use of pattern density control features. As already mentioned earlier, some fabrication processes, such as plasma etching and CMP, are affected by the local density of light or dark patterns: for example, the etch rate for silicon plasma etching depends on the local ratio between exposed silicon and photoresist, which can affect the resulting waveguide linewidth. A solution, also used widely in electronics design, is to use fillers or dummies to average the local density up or down to a standard value. This way, the layout designer has less to worry about how neighboring features will affect each other's behavior. Of course, this technique, often referred to as FILL, can only be applied when there is sufficient available space to add these patterns. Some photonic components, such as arrayed waveguide gratings, contain large uniform areas that can dramatically disturb the local average pattern density. This disruption in pattern density can affect photonic devices several hundred micrometers away from that area [32].

Similarly, the photonics designer is likely to benefit from optical proximity correction (OPC) modifications made to the as-designed layout so that the as-fabricated (on-wafer) geometries are as close to the layout as possible. Generally, OPC modifications are made by the foundry or mask-making house. Particularly for

very fine feature structures, these OPC modifications are important in order to reliably pattern features. The photonics design should be alert that in the case of fine features (e.g., sharp corners), in any event the fabricated structures will deviate in systematic ways from the "ideal" layout; lithographic TCAD tools can be used to better approximate the as-fabricated device structure if necessary.

10.5.5 Photonics design for test and design for reliability

This section has focused on existing and emerging methods for photonics DFM. As the level of complexity of photonic circuits continues to increase, development of new methods for photonics DFT will also become essential. Just as in electronic circuits, DFT will be necessary to assure not only designers but also customers that a circuit is operating correctly, as well as to debug circuit failures arising not only from design flaws, but also poorly understood interactions between manufacturing variations and the design. Thus, DFM and DFT methods in photonics are likely to be an important and coupled area for methodology development in the coming years.

For some photonic system designs, design for reliability may also be important. For the most part, passive photonic components do not suffer from wear-out or lifetime degradation effects, and reliability analysis is rarely necessary. A notable exception, however, are discrete or integrated lasers [39] and other active photonic device components. In systems incorporating active photonic devices, or high-power electronic components (e.g., resistive heating), careful analysis and design may be required to consider lifetime and usage requirements.

10.6 Future photonic-integrated circuit design flow requirements

As of this writing, there are still several areas that require additional development for the PIC design community. As PIC designs become more complex, the use of additional automation and design abstraction become desirable to enable designers to deal with greater complexity. Also, as integrated photonics reaches into different end applications the need for higher volume manufacturing and the ability to enable second-source manufacturing becomes important. In the electronic design flows, this was achieved through the use of automation tools like logic synthesis which was able to remap and reoptimize designs between different foundry libraries and standardized methodologies that enabled autoplacement and autorouting of synthesized designs. There is some automation for PIC physical design at present, mostly in the form of programmable cells and script-driven placements but work needs to be done to create a fully constraint-driven automated place and route flow for integrated photonics. Similarly, there are limited capabilities to map circuit-level components between foundries but there are no tools that automate the optimization of a design at the circuit level.

The integrated PIC industry is also starting to wake up to the benefits of circuit-level simulation, but adoption of circuit-level simulation has been slow, caused by a lack of complexity of designs in the transceiver space and due to the lack of simulation models backed by measured data from the foundries. This is exacerbated by the fact that foundries are still going through a steep learning curve to understand the nuances of photonic simulation modeling coupled with the lack of an industry standard modeling language to model photonic building blocks. While the PICs are not that complex, the systems into which they fit can be quite challenging given high-speed RF signals traces that must be carefully impedance-matched to get best performance. This highlights the need for better integrated parasitic extraction, thermal and stress analysis into the simulation flows.

Lastly, most of the CAD flow attention to date has been focused on design of photonic ICs, and very little attention has been given to DFM and DFT for integrated photonics. This includes subjects such as the use of lithography tools to review "as drawn" versus "as fabricated" structures. Finally, very little work has gone into the automation of integration of PICs with other ICs, whether photonic or electronic. We do see more demand in the industry for exploration of 2.5D and 3D semiconductor systems, including PIC or optical chiplets. And EDA vendors start to develop solutions to address the associated design challenges of bringing together ICs from different foundries, different domains (digital, analog, photonic), combined with interposers and packages.

References

[1] W. Bogaerts, L. Chrostowski, Silicon photonics circuit design: methods, tools and challenges, Laser Photon. Rev. 1700237 (2018) 1−29.
[2] M. Heins, et al., Design flow automation for silicon photonics: challenges, collaboration, and standardization, in: L. Pavesi, D.J. Lockwood (Eds.), Topics in Applied Physics, vol. 122, Springer, Berlin, Heidelberg, 2016, pp. 99−156.
[3] L. Chrostowski, M. Hochberg, Silicon photonics design, Self-published via Lulu.com, 2013.
[4] J.K. Patel, E. Ghilino, T. Korthorst, Are electrical circuit languages robust enough for photonics, Synopsys Tech. Pap., 2021.
[5] M.U. Khan, Y. Xing, Y. Ye, W. Bogaerts, Photonic integrated circuit design in a foundry + fabless ecosystem, IEEE J. Sel. Top. Quantum Electron. 25 (5) (2019) 1−14.
[6] Z. Lu, et al., Performance prediction for silicon photonics integrated circuits with layout-dependent correlated manufacturing variability, Opt. Express 25 (9) (2017) 9712−9733.
[7] W. Bogaerts, U. Khan, Y. Xing, Layout-aware yield prediction of photonic circuits, in: IEEE Int. Conf. Gr. IV Photonics, 2018, 93−94.
[8] J. Ding, P. Yang, M.I. Mishchenko, R.D. Nevels, Identify the limits of geometric optics ray tracing by numerically solving the vector Kirchhoff integral, Opt. Express 28 (7) (2020) 10670.
[9] B.M. Dillon, J.P. Webb, A comparison of formulations for the vector finite element analysis of waveguides, IEEE Trans. Microw. Theory Tech. 42 (2) (1994) 308−316.

[10] R. Scarmozzino, A. Gopinath, R. Pregla, S. Helfert, Numerical techniques for modeling guided-wave photonic devices, IEEE J. Sel. Top. Quantum Electron. 6 (1) (2000) 150−162.
[11] D. Gallagher, T. Felici, Eigenmode expansion methods for simulation of optical propagation in photonics: pros and cons, Integrated 4987 (2003) (2003) 69−82.
[12] W.P. Huang, C.L. Xu, S.-T.T. Chu, S.K.K. Chaudhuri, The finite-difference vector beam propagation method—analysis and assessment, J. Light. Technol. 10 (3) (1992) 295−305.
[13] K.S. Kunz, The Finite Difference Time Domain Method for Electromagnetics, CRC Press, London, UK, 1993.
[14] W.W. Chow, S.W. Koch, M. Sargent, Semiconductor-Laser Physics, Springer, Berlin/Heidelberg, 1994.
[15] A. Rahim, T. Spuesens, R. Baets, W. Bogaerts, Open-access silicon photonics: current status and emerging initiatives, Proc. IEEE 106 (12) (2018) 2313−2330.
[16] J. Pond, X. Wang, Z. Lu, E. Schelew, G. Lamant, A. Farsaei, Latest advancements to the industry-leading EPDA design flow for silicon photonics: invited paper, in: 2019 IEEE/ACM International Conference on Computer-Aided Design (ICCAD), 2019, pp. 1−6.
[17] P. Gunupudi, T. Smy, J. Klein, Z.J. Jakubczyk, Self-consistent simulation of opto-electronic circuits using a modified nodal analysis formulation, IEEE Trans. Adv. Packag. 33 (4) (2010) 979−993.
[18] M. Fiers, et al., Time-domain and frequency-domain modeling of nonlinear optical components at the circuit-level using a node-based approach, J. Opt. Soc. Am. B 29 (5) (2012) 896.
[19] C. Arellano, S.F. Mingaleev, E. Sokolov, I. Koltchanov, A. Richter, Time-and-frequency-domain modeling (TFDM) of hybrid photonic integrated circuits, in SPIE OPTO, vol. 8265, 2012, 82650K.
[20] A. Farsaei, et al., A novel and scalable design methodology for the simulation of photonic integrated circuits, in: Adv. Photonics 2016 (IPR, NOMA, Sensors, Networks, SPPCom, SOF), vol. 2016, 2016, 3−5.
[21] J. Patel, J. Patel, et al., Electro-optic & custom photonic co-design in monolithic silicon photonics technology, in: DesignCon 2022 Proceedings.
[22] W. Bogaerts, M. Fiers, M. Sivilotti, P. Dumon, The IPKISS photonic design framework, in: Optical Fiber Commun. Conf. 2016, 2016, pp. 2−4.
[23] R. Cao, J. Ferguson, F. Gays, Y. Drissi, A. Arriordaz, I. O'Connor, Silicon photonics design rule checking: application of a programmable modeling engine for non-Manhattan geometry verification, in: IEEE/IFIP International Conference on VLSI and System-on-Chip, VLSI-SoC, vol. 2015, 2015, pp. 1−6.
[24] R. Cao, J. Ferguson, A. Bakker, T. Korthorst, A. Arriordaz, I. O'Connor, Photonic designs with EDA approach: a novel methodology for curve design validation, in: 2015 IEEE Summer Topicals Meeting Series, SUM 2015, 2015, pp. 29−30.
[25] R. Cao, et al., LVS check for photonic integrated circuits − curvilinear feature extraction and validation, in: Proc. of the 2015 Design, Automation & Test in Europe Conference & Exhibition, 2015, pp. 1253−1256.
[26] H. Bahrami, H. Sepehrian, C.S. Park, L.A. Rusch, W. Shi, Time-domain large-signal modeling of traveling-wave modulators on SOI, J. Light. Technol. 34 (11) (2016) 2812−2823.
[27] B. Wang, et al., A compact verilog—a model of silicon carrier-injection ring modulators for optical interconnect transceiver circuit design, J. Light. Technol. 34 (12) (2016) 2996−3005.

[28] V. Stojanović, et al., Monolithic silicon-photonic platforms in state-of-the-art CMOS SOI processes [invited], Opt. Express 26 (10) (2018) 13106.
[29] Y. Xing, J. Dong, U. Khan, W. Bogaerts, Hierarchical model for spatial variations of integrated photonics, in: IEEE Int. Conf. Gr. IV Photonics, 2018, 91−92.
[30] K. Qian, B. Nikolić, C.J. Spanos, Hierarchical modeling of spatial variability with a 45nm example, in: Des. Manuf. Des. Integr. III, vol. 7275, 2009, 727505.
[31] B.E. Stine, D.S. Boning, J.E. Chung, Analysis and decomposition of spatial variation in integrated circuit processes and devices, IEEE Trans. Semicond. Manuf. 10 (1) (1997) 24−41.
[32] Y. Xing, J. Dong, U. Khan, W. Bogaerts, Correlation between pattern density and linewidth variation in silicon photonics waveguides, Opt. Express 28 (6) (2020) 7961−7968.
[33] B.E. Stine, et al., Rapid characterization and modeling of pattern-dependent variation in chemical-mechanical polishing, IEEE Trans. Semicond. Manuf. 11 (1) (1998) 129−140.
[34] D.S. Boning, S.I. El-Henawy, Z. Zhang, Process variation-aware photonic design, in: Optical Fiber Communication Conference (OFC)2, 2021, p. Tu5B.1.
[35] S.I. El-Henawy, C. Lang, D.S. Boning, Yield prediction for coupled-resonator optical waveguides using variation-aware compact models, in: Frontiers in Optics, 2019, p. FTu1C.3.
[36] D. Boning, S.I. El-Henawy, Z. Zhang, Variation-aware methods and models for silicon photonic design-for-manufacturability, J. Light. Technol. (2021) 1.
[37] Z. Zhang, et al., Enabling wavelength-dependent adjoint-based methods for process variation sensitivity analysis in silicon photonics, J. Light. Technol. 39 (6) (2021) 1762−1769.
[38] R. Biernacki, S. Chen, G. Estep, J. Rousset, J. Sifri, Statistical analysis and yield optimization in practical RF and microwave designs, in: 2012 IEEE/MTT-S International Microwave Symposium Digest, 2012, pp. 1−3.
[39] J.C. Norman, et al., Reliability of lasers on silicon substrates for silicon photonics, Reliability of Semiconductor Lasers and Optoelectronic Devices, Elsevier, 2021, pp. 239−271.

This page intentionally left blank

Photonic-integrated circuit fabrication and test approaches

11

Ajey P. Jacob[1], *Sujith Chandran*[1], *Joris Van Campenhout*[2], *Marianna Pantouvaki*[2], *John Heck*[3], *Ken Giewont*[4], *Michal Rakowski*[4], *Yusheng Bian*[4], *Peter Debackere*[5] *and Matthias Kuntz*[5]

[1]Information Sciences Institute, University of Southern California, Marina Del Rey, CA, United States, [2]imec, Leuven, Belgium, [3]Intel Corporation, Santa Clara, CA, United States, [4]GlobalFoundries, Malta, NY, United States, [5]Infinera, Sunnyvale, CA, United States

11.1 Introduction

Photonic-integrated circuits (PICs) integrate electronics and photonics elements on a silicon, indium phosphide (InP), or a heterogeneous material platform for various applications, such as data centers, communications, LIDAR, sensors, and spectroscopy, besides several defense application sectors. An advantage of InP-based PICs and its compounds indium gallium arsenide phosphide (InGaAsP), and indium gallium aluminum arsenide (InGaAlAs) is that the material system supports light generation, amplification, modulation, and detection, integrating all required functions for sophisticated optical transmitters and receivers, making monolithic integration possible for a wide range of functionalities. However, silicon-based PICs have democratized the technology among the various platforms due to their ability to leverage microelectronics manufacturing infrastructures, such as design, validation, fabrication, characterization, and packaging. In addition, extensive research in the field since the early 2000s has enabled silicon photonics high volume manufacturing in advanced foundry ecosystems to create energy-efficient scaled photonics systems. Today, both silicon-based and InP-based PIC platforms are manufactured and sold for several applications, including long-range and short-range communications. This chapter discusses silicon and InP-based photonics chip integration and fabrication schemes, and testing.

11.2 Relevant industry trends

Data generation has seen unprecedented growth in recent years, driven by the mobile and internet of things (IoT) sectors. It has resulted in a surge in demand for the PIC, primarily implementing long- and short-distance communication networking and cloud infrastructure. For example: by 2025, long-haul telecom traffic is

Integrated Photonics for Data Communication Applications. DOI: https://doi.org/10.1016/B978-0-323-91224-2.00012-6
© 2023 Elsevier Inc. All rights reserved.

expected to grow by over 42% and cloud data centers by 30%, increasing data traffic from 400 Gbps to 1 Tbps. While the intradata centers focus on pluggable standards beyond 400 G, the regional and metro interdata centers push for spectral efficiency through higher-order modulation. Furthermore, as the solid photonic interconnect market grows, the 5- and 6-G radio access networks will also benefit through their synergy with data center optically compatible architectures and protocols.

11.3 Photonic-integrated circuit fabrication schemes and methodology

A PIC for transceivers consists of (1) laser light sources, (2) passive and active devices that convert analog light signals to digital electronic data, (3) driver and amplifier electronics circuits, and (4) application-specific ICs. In response, the industry has developed multiple PIC integration schemes for silicon photonics technology. These can be categorized into four main types: (1) hybrid, (2) heterogeneous, (3) monolithic, and (4) fully monolithic. Fig. 11.1 shows a schematic illustration of the integration schemes. The III−V photonics also follow similar integration schemes. In a hybrid PIC, the laser, photonics, and electronics components are separately fabricated and then packaged. In a heterogeneous PIC, the laser and photonics IC is fabricated together, and the electronics IC is then packaged. In a monolithic PIC, the photonics and electronics components are fabricated together, and the laser is then packaged. Finally, in a fully monolithic PIC, laser, photonics, and the electronics components are fabricated together with minimal packaging providing high bandwidth density and high energy efficiency.

The following section shall discuss PIC fabrication through a (1) brief discussion on the design methodology that leads to the PIC fabrication, (2) detailed discussion on various PIC integration schemes, (3) brief discussion on PIC packaging, and (4) brief discussion on test methodology. Successful implementation of the fabrication methodology leads to improvements in performance, throughput, yield, and cost of a PIC. In addition, this standardized methodology provides freedom and flexibility to the designer to submit a custom photonic device layout for fabrication.

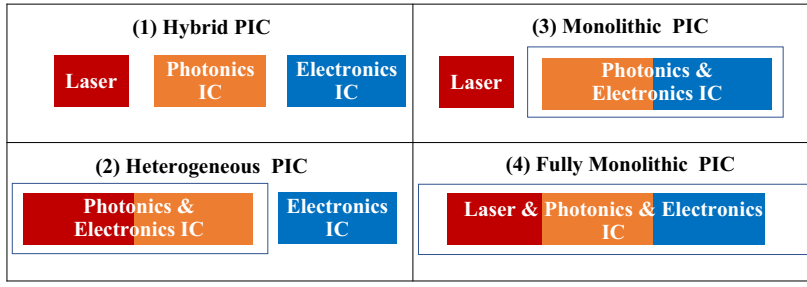

Figure 11.1 Illustration of four types of silicon photonics integration schemes.

Following the PIC fabrication, a system can be built through the cointegration of multiple components implemented at the chip level, wafer level, or interposer level using flip-chip bonding, through silicon vias (TSVs), or other advanced optoelectronics packaging schemes.

11.3.1 Photonic-integrated circuit physical design flow

PIC high-volume manufacturing industry follows a similar systematic logistical approach as in a CMOS foundry with a series of established protocols assisted by computer-aided design and electronic design automation (EDA) tools in translating an abstract idea into a reproducible fabrication line. EDA tools have supported the design and evaluation of built-in circuits and components, physical design flow, and the statistical analysis of each element to monitor process variations, resulting in high throughput, yield, and cost-effectiveness. Fig. 11.2 illustrates a PIC manufacturing logistics flow. The set logistics follow the photonics foundry PDK that contains a library of the basic photonics and electronics building blocks for the integrated circuits, along with a set of design rules, as shown in Fig. 11.3 and as discussed in detail in Chapter 10. By using the PDK, photonics components are designed and simulated using physics-based equations, such as Maxwell equations. The device response is then translated into a compact model (transfer function) in order to perform circuit-level simulations to elucidate the system-level functionalities of interconnected components in a timely and repetitive manner. To prepare the PIC design for foundry tape out, physical design flow, such as layout generation, routing, verification, design rule check (DRC), and layout versus schematic (LVS), is followed. Layout is a set of mask layers for patterning the devices during fabrication.

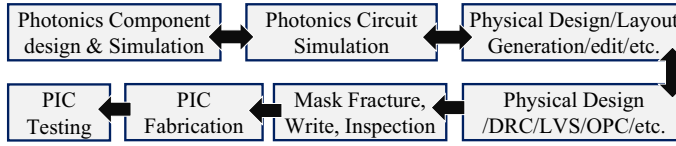

Figure 11.2 Photonic-integrated circuit manufacturing logistics.

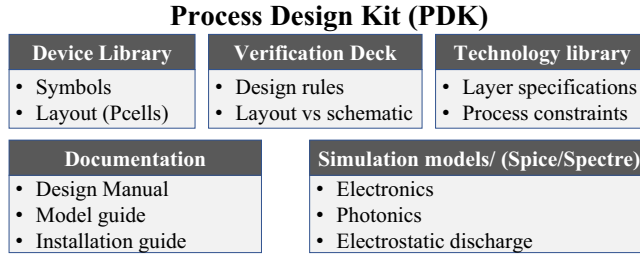

Figure 11.3 Typical process design kit components for photonic-integrated circuit manufacturing.

The layout must be verified against the design rules because of lithography's process limitations (e.g., minimum feature size, spacing, area). Implementing these design rules ensures reproducible performance and maximum yield. When a layout meets the design rules, it is taped out. The fabricated device is then characterized, and the hardware response is compared to the simulation results for maximum yield. In most cases, the initial design contributes to the ideal hardware since the design assumptions are based on the respective technology's process design kit (PDK).

11.3.2 Device design and simulation flow

Following initial device design built on underlying physics-based theoretical considerations, the device geometry and its responses using well-evaluated material models available in photonics design software libraries are verified. Commonly used computation techniques for photonics devices are finite difference time domain (FDTD), beam propagation method or finite difference eigenmode solvers [1–3]. The beam propagation method approach is based on slowly varying envelope approximation, and it is used for evaluating the light propagation in low index contrast waveguide structures [4]. FDTD directly solves Maxwell's equations in the time domain, and it is one of the most accurate numerical methods for light propagation in photonic device design [5]. However, it is computationally intensive and time-consuming. Additionally, its simulation accuracy strongly depends on the simulation grid size. Accurate simulations demand a smaller simulation grid and thus a high-performance computing infrastructure. In contrast, mode solvers effectively estimate waveguide properties, such as effective index, group index, and dispersion parameters [1,6]. This simulation tool is used for optimizing the waveguide geometry for single mode guiding conditions. A mode solver is also found to be suitable for the semianalytic modeling of longer structures like high-quality factor microring resonators, Mach–Zehnder interferometer (MZI) [7], and arrayed waveguide gratings (AWGs) [8].

To optimize the device performance, multiple numerical simulations are carried out by changing critical geometrical parameters. However, for large-scale integrated circuits on a 300 mm wafer, device layer thickness and process conditions can be nonuniform from chip to chip, wafer to wafer, and lot to lot. Considering this uncertainty, photonics device geometry must be optimized to provide a suitable response in the presence of potential variations in the material and process.

Once the optical circuit simulations are completed using time or frequency domain simulations, the device response corresponding to the device input and output is implemented using a scattering matrix or with empirical/analytical models or compact models, for example in Verilog-A, an industry standard modeling language. Such simulation blocks (schematic) can then be further used for circuit-level verifications.

11.3.3 Layout generation

The photonics design flow includes layout generation. It is imperative to note that rectilinear electronic layouts differ significantly from curvilinear photonic layouts.

Manufacturing design flows that follow electronic design rules (rectilinear configuration) cause significant problems during the photonics circuit DRC. As an example, almost all passive photonic structures, such as straight waveguide, bend waveguide, splitters/couplers, and filters, can be defined with one or two mask layers. However, for monolithic integration schemes in which silicon electronics and silicon photonics devices are fabricated sequentially, critical steps, such as implantation and metallization, can adversely affect the performance of the photonics devices. Consequently, a few mask layers are added to photonics layers to protect the waveguide from the metal and implant contaminants. Often the electronic device layout contains parameterized cells (p-cell) to scale the device. However, for photonics devices, parameterization is usually limited only to the length of the waveguide, as scaling on other parameters can significantly change the device response. To avoid potential errors in the photonic component layout design, simulation layout (from FDTD or beam propagation methods) for the waveguide layer (mask layers define the waveguide cores) can directly export as a p-cell, and most simulation software support this feature. This option also allows importing the mask layout (generally GDS II format) to the simulation window for verification (if the layout is modified to satisfy design rules).

11.3.4 Design rule check

Design rules are a set of constraints imposed on the geometry of the layout to ensure maximum yield for the components of the PIC in high-volume manufacturing. Just as in electronic-integrated circuits, the rules are created for a particular technology based on the fabrication process to reproduce the intended design. A DRC is a process for checking (physical verification signoff) that the layout complies with the specified geometrical constraints. This usually involves checking minimum feature size, minimum separation and enclosed area, acute angle, missing layer, misalignment, notch spacing, and pattern density. In the absence of Manhattanized curvilinear design rules, rectilinear design rules can result in extremely time-consuming and tedious DRC checks (see section for Manhattanization for more details). Another important aspect of physical verification is the LVS. LVS determines if the layout and schematic defined by the circuit are identical.

11.3.5 Mask rule check

Mask rule checking (MRC) ensures that manufacturing limits for masks are being followed and that they are being represented correctly before printing. This rule is designed to help prevent the need for costly and time-consuming rewrites of masks by catching errors before they are written. As far as rectilinear structures are concerned, the MRC has been well developed; however, a reliable set of MRC rules for curvilinear structures needs to be developed. As a result, curvilinear shapes cannot be checked for width and space using the simple set of rules that apply to rectilinear shapes in order to determine whether they are wide or small. There is a critical role for the MRC when curvilinear structures require OPC to use them in

lithography processes. Every mask process, for example, has resolution limits for exposed and unexposed features and two-dimensional limits for corner-to-corner distances. As a result, OPC engines are now configured to allow mask manufacturing limits to be set to avoid violations. It can also so happen that after OPC, the standard mask rule limits may also be violated. Although there are similar conditions for making masks for rectilinear and curvilinear structures, the rules' limits need to be revised despite the fact that the requirements for mask making are identical. The set of MRC rules include (1) checks on internal and external dimensions, (2) checks on minimum area, and (3) checks on curved contour shapes, such as convexity and concavity. In internal and external dimension checks, the distance between edges is measured either from inside the exposed feature (internal) or outside the next exposed feature (external). As the minimum area checks the minimum distance from the edges, the maximum curvature checks are a supplementary type of MRC specific to masks with curvilinear patterns [9,10].

11.3.6 Design rule challenge: Manhattanization of curvilinear structures

For photonics circuits, curvilinear shapes facilitate a smooth path for light propagation and satisfy the principle of total internal reflection. Conventional CMOS physical design flow includes DRC, optical proximity correction (OPC), MRC, mask process correction (MPC), mask writing, and mask inspection. These fabrication tape-out steps work well for orthogonal shapes with vertical and horizontal edges, but they often perform poorly for curvilinear shapes. One approach to solve these problems is to convert curvilinear edges to small vertical and horizontal fragments that closely match the design—a process known as Manhattanization [11,12]. With this additional step, specific ground rules (DRC, OPC, MRC, and MPC) defined for rectilinear structures can be applied to the curvilinear photonic structures. However, this can cause problems, such as the need for fine Manhattanization grids, especially for structures with small radius of curvature. This results in fragments with feature size less than what the technology can support. If the Manhattanization grid size is increased to account for minimum feature size, the Manhattanized curvilinear structures turn out to be quite different from the intended shape. Fig. 11.4A shows the layout of a rectangle rotated at an arbitrary angle according to the design requirement. The design rules are based on orthogonal edges; tilted structures (rectangles at arbitrary angles) must be Manhattanized. As an example, Fig. 11.4B shows the Manhattanized version of Fig. 11.4A with a 10-nm grid size. As expected, the Manhattanized layout contains notches of 10-nm wide (see Fig. 11.4B) at the corners of the rectangle, and these notches cause design rule violations as the width of the notches is less than the minimum feature size offered by the technology. To overcome the problem associated with the minimum feature size, additional features must be added to the layout regions of violation (see Fig. 11.4C). Comparing Fig. 11.4A and B, it is evident that the DRC clean layout is entirely different from the designer's intention. Fig. 11.4D shows the layout view of a second order

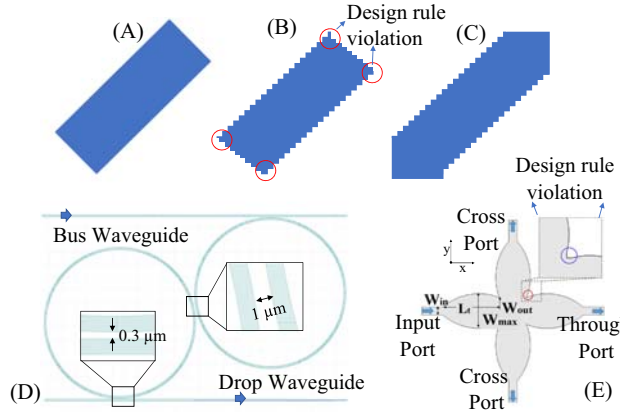

Figure 11.4 (A) Designer intended layout, (B) Manhattanized layout, and (C) design rule check clean layout. (D) Layout view of a ring resonator-based second order wavelength filter; inset shows the magnified views of the coupler sections with different waveguide separation. (E) Compact waveguide crossing with curvilinear shapes; inset shows design modification to satisfy the design rules.

microring filter with coupler sections between the rings and between straight and ring waveguides [13]. The performance of the filter strongly depends on the coupler section, and any geometrical (width, height, waveguide separation, and sidewall angle) variation can significantly change the filter response. Magnified views of the coupler sections (straight to bend and bend to bend) are shown as a figure inset. Note that the waveguide separation for the straight to bend coupler and bend to bend coupler are different for the same waveguide width. This is a concern for lithography (due to proximity) and etching (due to loading effect). Generally, a structural offset (in width) is observed between postlithography and postetch structures. This offset can be corrected with the OPC flow for the isolated photonics structures. However, this may be a critical issue with the coupled waveguide structures with varying pitches (separation). For a photonic waveguide, width offset at the top, middle, and bottom along the thickness has varying influences on the performance. A waveguide designed to operate for transverse electric (TE) polarization is sensitive to the structural variations in the mid-height of the waveguide. Consequently, the OPC is modified to match the mid-height width of the waveguide to the design. However, this remains a critical problem as the waveguide cross-section profiles, and sidewall angles can be significantly different for isolated, nested, and coupled structures. So, mass correction of the etch profile turns out to be exceedingly difficult. Fig. 11.4E shows the schematic top view of a compact (4.7×4.7 μm) waveguide crossing [14]. An ideal crossing structure includes acute angles and zero-gap regions at the corners of the crossing junction. This invokes ground-rule violations in terms of minimum feature size and acute angle. The inset figure shows a modified design of the crossing that satisfies ground rules at the expense of the device performance. To ensure

the maximum performance of the photonics structures, existing ground/design rules are modified to fit the curvilinear shapes.

Manhattanization with a larger grid size can significantly change the intentional design layout, whereas Manhattanization with smaller grid size can invoke significant design rule violations. Process variation further amplifies the difference between the design and hardware. In photonics, slight variations in the geometry can induce phase error due to the high refractive index property. This will have a negative impact on resonant and interferometric structures. Thus mask rules for curvilinear structures will eliminate the problems caused by Manhattanization. Mask rules for the curvilinear structure must consider features associated to its minimum width, space, area, and acute angles.

After the wafers are fabricated, facilitating wafer acceptance through effective testing for high precision, repeatable, and reliable wafers is an essential aspect of manufacturing. The following section will discuss in-line and end-of-line wafer testing.

11.4 Photonic-integrated circuit in-line and end-of-line testing

As in any integrated chip manufacturing, PICs undergo testing during (in-line) and after their fabrication (end-of-line) to quantify the statistical variation and yield losses. In addition, optoelectronic test structures are created for photonic devices to test their performance, just as in an electronic IC scribe line structure. However, testing integrated photonic devices is much more complex than electronics as the latter only requires a firm electrical contact within the metal pads of tens of square micrometers. In contrast, photonic device characterization demands precise nano-scale fiber alignment to couple light in and out of the photonics chips.

Identifying early failures and their causes requires high-throughput automated photonic testing. The commonly used parameters to identify early failures, such as optical loss, are wavelength, frequency, voltage, temperature, and optical power. Wafer-level inline and chip-level end-of-line optical characterizations are the two testing approaches conducted on photonics devices. In-line wafer-level testing is performed by using grating coupling [15], end-of-line testing utilizes either grating coupling or butt (end-fire) coupling mechanisms. Butt-coupling uses a pair of lensed fibers or lenses with high magnification to focus light on the waveguide's input and output facets [16,17]. The grating coupler testing approach offers flexibility and is tolerant to alignment inaccuracy; however, its shortcomings are wavelength-dependent throughput, high insertion loss, and polarization dependency. Butt coupling, on the other hand, provides wavelength and polarization independent performance and lower insertion loss—mode-mismatch and Fresnel reflections contribute to optical loss. The EOL approach may require preprocessing of the chips, such as dicing and polishing. It also demands stringent alignment of the characterization setup. Inline testing primarily qualifies component-level design

accuracy and process-related variations, while EOL testing accurately quantifies statistical production yield. Generally, die-to-die, die-to-wafer, and wafer-to-wafer testing and correlation are studied.

Fig. 11.5 shows photographs of an automatic in-line photonic characterization system from Formfactor. These tester units can provide high throughput, cost-efficient wafer-level electro-optic tests through automated scripts to extract model parameters in a highly efficient manner. This can also be employed for the detailed characterization of the passive devices to aid their process development and optimization. The automated tester consists of a custom photonic test head connected to a standard 300 mm wafer prober. Vertical grating couplers and fiber optic cables bring C-, L-, or O-band optical signals to and from a wafer while a standard electrical probe card handles electrical signals. Typical measurements for the passive photonic devices (Reference: Chapter 5) include propagation loss, wavelength-dependent insertion loss, polarization-dependent insertion loss, crosstalk, and return loss. Interferometric structures (MZI or microring resonator) are used for the extraction of waveguide properties, such as effective index and group index (Reference: Chapter 3). Photodetectors (PD) are characterized in terms of dark current, responsivity, and bandwidth (Reference: Chapter 4). Electro-optic modulators are tested for phase efficiency/shift, insertion loss, rise time, and fall time (Reference: Chapter 3). From the test data, one can extract $V_{\pi}L$, an essential metric for modulators, where V_{π} is the voltage required for π phase shift and L is the length of the modulator. A subset of these characteristics is sampled on production wafers to monitor the process health and identify outlier wafers unsuitable for shipment for laser and gain element integration. Due to the grating coupler efficiency variability, most optical parameters of interest cannot be measured directly. Instead, a cut-back testing method can be used to derive the underlying parameter from a series of related test structures [18]. Linear interpolation on the results eliminates losses due to the grating couplers and splitters and calculates the average loss per length of the waveguide. The automatic testers align fibers to the gratings with an accuracy and repeatability better than 0.4 dB (1-sigma).

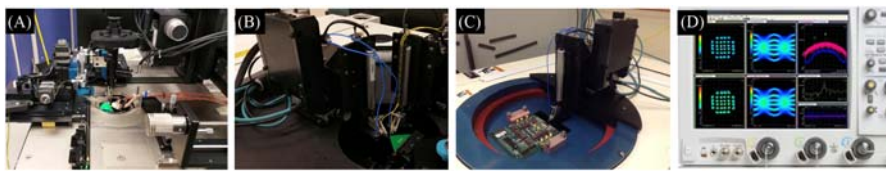

Figure 11.5 Photographs of a commercial 300-mm wafer probe platform with RF and photonic capability, including automated multifiber alignment for C- and O-band testing. (A) Wafer characterization tool with 67-GHz light wave tunable laser and optical vector analyzer, (B) inline optical test tool with low current DC measurements, (C) wafer product tester with dc and analog electrical test capability, (D) module characterization with high-speed digital and optical functional test capability for continuous wave and module optical evaluation module [19].

11.5 Photonic-integrated circuit integration, fabrication, and manufacturing

The following sections will discuss in detail various integration schemes, (1) hybrid, (2) heterogeneous, (3) monolithic, and (4) InP Photonics, used for PIC fabrication and manufacturing.

11.5.1 Hybrid photonics integration platform

Hybrid photonics is a mature integration technology where the critical functional blocks, such as III−V lasers, silicon drivers, amplifiers, and ASICs, are fabricated separately and packaged together in a flip-chip multichip package or on a silicon optical interposer. While the silicon photonics components are fabricated on a silicon-on-insulator (SOI) platform, the electronic components are built on bulk silicon. As silicon-based efficient optical sources are fundamentally challenging because of silicon's indirect bandgap, III−V-based lasers, such as VCSEL, are independently manufactured on the III−V line, which also avoids III−V contamination of the silicon manufacturing line. They are then packaged with the silicon PIC. Hybrid integration schemes give flexibility in the choice of lasers. However, the link budget associated with the lasers will be extremely high because of the coupling loss. Therefore the hybrid lasers on the PIC are evanescently coupled. While the ASICs are fabricated on advanced silicon processes, such as those below 14-nm finFET, there is flexibility in manufacturing the silicon photonics components on older process nodes, such as 180, 130, or 90 nm. The lagging photonics process nodes may be manufactured with a 200-mm dry lithography tool or an advanced 300-mm immersion lithography tool. Thus the hybrid silicon platform combines the maturity and scaling of CMOS processing and SOI photonics with essential III−V active components, such as lasers. Most of the leading silicon photonics foundries follow this integration scheme, while the III−V InP photonic platform codesigns with silicon germanium (SiGe) electronics drivers.

This section will give an overview of silicon photonics fabrication, and in particular an interposer platform, enabling hybrid integration of electro-optical modules with TSV-assisted electrical connectivity to a packaging substrate. The silicon photonics interposer process flow is described first, followed by a discussion of the electrical and optical characteristics of some key building blocks that form a proof-of-concept hybrid EIC−PIC transceiver.

11.5.1.1 Design considerations

Hybrid integration is different from monolithic integration in that it allows for combining individually optimized PIC and EIC technologies. For example, advanced CMOS FinFET nodes can be employed to realize power-efficient, high-speed modulator driver and transimpedance amplifier (TIA) circuits, as well as other digital circuitry. Flip-chip assembly using copper (Cu) microbumps can be used to electrically connect the EIC to the PIC, with interconnect pitches down to 50 μm and

parasitic capacitance of just a few tens of femto Farads. For connectivity to the substrate package, (commercial) demonstrations of hybrid integrated optical modules have so far mostly relied on wire bond interconnects. However, as transceivers scale in the transmission rate and capacity in each lane, it is increasingly challenging to retain signal integrity of the high-speed RF signals, as well as to manage power delivery to the module, by just using wire bonds. The adoption of TSVs can help to mitigate these bottlenecks. To this end, TSVs could be integrated either in the EIC or the PIC. Implementation of TSVs in the PIC allows the PIC to be implemented as the bottom chip, with the EIC placed on top [20]. This configuration keeps the top surface of the PIC accessible for optical coupling to fiber and the EIC, and potentially a hybrid integrated laser, in close proximity to a heat sink placed on top of the module.

TSVs can be integrated in silicon photonics platform either in the middle of line, before the back end of line (BEOL) modules, or at the end of line. The former approach typically enables more aggressive pitch scaling and high RF bandwidth and is as such more future proof. However, a TSV-middle integration approach typically imposes a limitation on the maximum interposer thickness. The discussion in this section will be based on IMEC's silicon photonics interposer platform, which contains middle-of-line TSVs with 10 μm diameter and 100 μm interposer thickness, which enables TSV pitch down to 40 μm while leaving sufficient margin for deeply etched structures in the PIC substrate, as required for optical edge coupling or local thermal isolation. In addition to the TSVs, passive optical waveguides and active photonic devices are cointegrated on a single SOI substrate using CMOS compatible front end of line (FEOL) modules, tungsten contacts, and copper interconnects. The platform uses 300-mm SOI wafers with a 220-nm crystalline silicon on a 2-μm buried oxide (BOX) and has the following features: (1) passive Si waveguides with three different etch depths (220, 150, and 65 nm), (2) 160-nm poly-Si layer for advanced passive optical devices, including grating couplers, (3) carrier depletion-based p−n junction Si ring modulators (RM), (4) metal heaters for thermal control, (5) germanium (Ge) p−i−n PD, (6) two-level Cu interconnects, (7) Cu through-silicon vias (TSV middle), (8) aluminum copper (AlCu) bond pads, (9) frontside Cu microbumps, and (10) backside Cu redistribution layer (RDL) and Cu microbumps.

Fig. 11.6 illustrates the cross-section of a hybrid silicon photonics platform with TSV integrated interposer technology. After the FEOL processing modules define the passive and active optical devices on the interposer, TSVs are fabricated for electrical connectivity to the backside of the wafers. At the same time, Cu microbumps on the front side enable flip-chip integration of the EIC die on the silicon photonics die. A TSV size of 10 μm diameter on a 100 μm Si interposer thickness is selected to ensure mechanical stability of the silicon photonics die (chiplet) during subsequent packaging after die singulation.

Fig. 11.7 illustrates silicon photonics interposer flow to describe the hybrid silicon photonics integration scheme. IMEC's silicon photonics interposer flow is described in a block diagram. The FEOL part of the flow consists of fabricating passive waveguides, Si modulators, and Ge PD. After that, the tungsten contacts

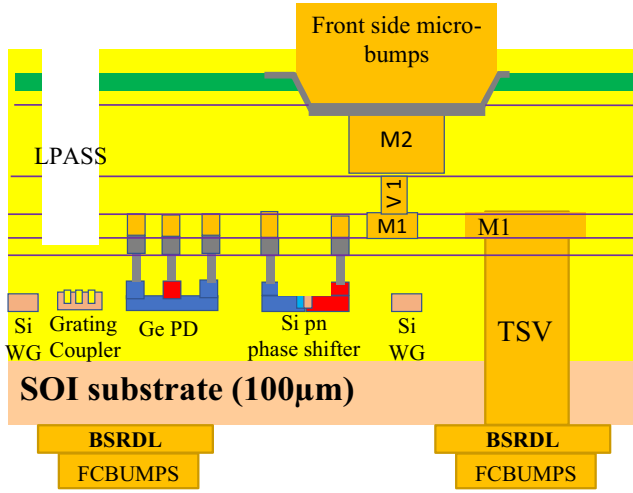

Figure 11.6 High-speed eye diagram generated by the CMOS-SiPh transceiver prototype: (A) Receiver electrical eye diagram at 40Gb/s NRZ measured at the TIA output of the receiver for 1339 nm wavelength, (B) transmitter optical eye diagram at 40 Gb/s NRZ for 1333.2 nm wavelength. The ER is 4.6 dB, and (C) TX to RX loopback eye diagram measured at 1333.2 nm wavelength for 32 Gb/s NRZ [21].

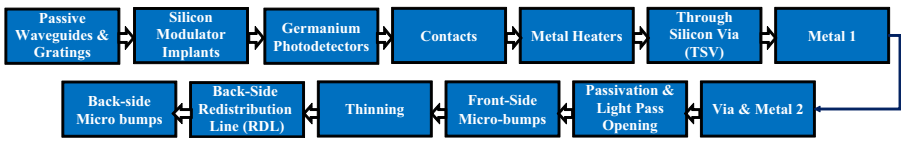

Figure 11.7 Block diagram illustrating the integration flow of imec's 300-mm hybrid silicon photonics interposer platform.

and heaters are fabricated. Finally, the TSV module is integrated following the TSV-middle approach, after the contacts and metal heaters and before the BEOL Metal 1, Cu interconnect layer, as shown in Fig. 11.6.

The TSV module includes the following process steps: silicon nitride (SiN) deposition, lithography, TSV etch and strip, liner deposition, barrier and seed deposition, TSV Cu-plating, anneal, and chemical mechanical polishing. The integration flow of the TSV module is shown in Fig. 11.8 [16]. TSV patterning is performed using 248-nm lithography. It involves etching the thick dielectric stack, including the 2-µm BOX of the SOI wafers and the TSV cavities, using a Bosch deep reactive ion etching process. An oxide liner is deposited after TSV etch to ensure electrical isolation of the TSVs from the Si substrate, followed by a pulsed vapor deposition (PVD) Tantalum (Ta) barrier layer and a PVD Cu seed required prior to Cu plating. Based on the conformality of the dielectric liner, the process needs to be adapted to ensure good coverage across the full TSV depth. Finally, TSV filling is performed by bottom-up electrochemical deposition Cu plating.

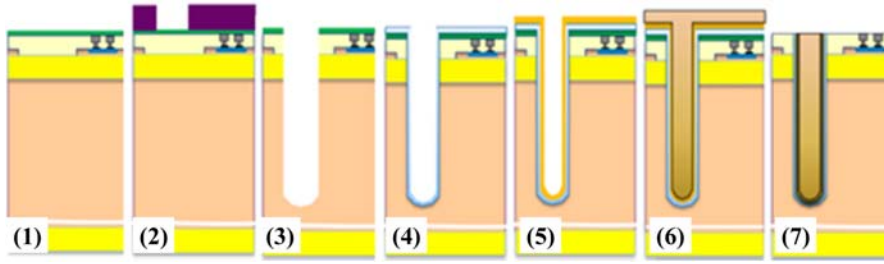

Figure 11.8 The major steps for through silicon via-middle integration scheme; (1) SiN deposition (2) Patterning (3) through silicon via/etch/strip, (4) O3/TEOS Liner deposition (5) barrier/seed deposition (6) Cu plating, and (7) Cu/Ta + liner/SiN CMP.

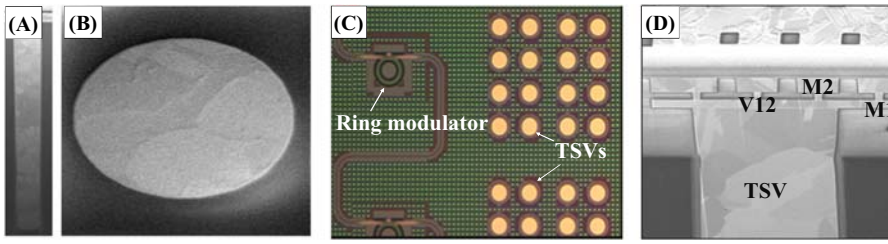

Figure 11.9 (A) Through silicon via cross-section after plating, (B) through silicon via after CMP, (C) microscope image after the through silicon via module, showing the cointegration of Si-ring modulators with through silicon vias, and (D) through silicon via cross-section with up to metal-2 interconnects.

In the TSV-middle approach, the TSVs are exposed to the thermal budget of the BEOL processing of the following modules (~400°C). A postplating anneal is commonly applied before Cu CMP to reduce the "Cu pumping" during BEOL processing [17]. High tensile stress across the wafer, caused by a mismatch between the coefficient of thermal expansion of Si and Cu, may generate strain in the Si around a TSV. Therefore a keep-out zone (KOZ) must be defined around the TSV to avoid the impact of this effect on the devices. The TSV module consists of Cu, Ta barrier, and dielectric liner/SiN CMP. The CMP step is critical to expose the heater metal to enable electrical contact with the BEOL layers. Therefore it is crucial to avoid any voids. After optimizing the plating process, no voids were observed after TSV CMP on the SOI wafers Fig. 11.9A and B. A microscope image from a silicon photonics interposer after completing the TSV module is shown in Fig. 11.9C and its cross-section image is shown in Fig. 11.9D.

After completing the front-side processing, the wafer is mounted on a carrier wafer with the backside facing up. The process flow for backside processing and respective SEM images taken at distinct stages of this flow are shown in Fig. 11.10. First, wafer thinning to 100 μm target thickness is achieved using a combination of coarse grinding and more refined polishing of the last few μm to reduce surface roughness.

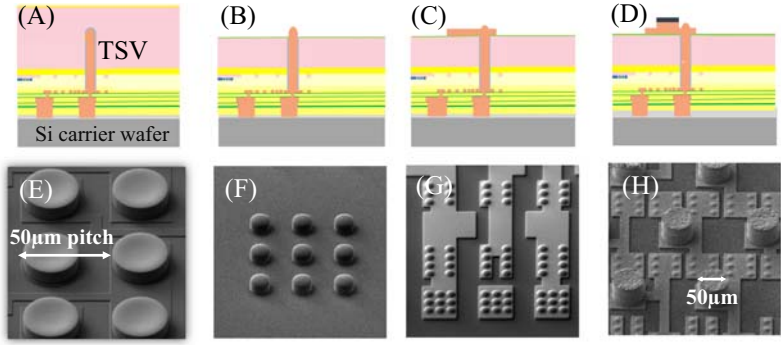

Figure 11.10 Schematic cross-sectional views of the back-side processing steps and the respective SEM images: (A) mounting on carrier wafer, (B) thinning and through silicon via reveal, (C) back-side redistribution layer, and (D) back-side microbumps. SEM images of (E) front-side microbumps, (F) back-side through silicon via reveal, (G) back-side Cu redistribution layer, and (H) back-side microbumps.

This is followed by exposure of the TSVs at the backside by Si etch-back, deposition of a SiN passivation layer, and TSV reveal (Fig. 11.10B). The next module aims at electro-plating the Cu RDL to interconnect the TSVs at the backside (Fig. 11.10C). Finally, backside microbumps are electro-plated for assembly on a PCB (Fig. 11.10D). For example, in [21], Cu/Ni/Sn pillars of 50 μm diameter and 150 μm pitch were used to assemble the SiPh interposer on a PCB using thermo-compression bonding. After backside processing, the wafer is debonded from the carrier wafer with the front side up to expose optical coupling through the front side of the SiPh interposer. After dicing, the SiPh is assembled to the package substrate by thermo-compression bonding. The CMOS die is then flip chip bonded on the SiPh interposer.

11.5.1.2 Hardware performance

A set of electrical and optical measurements are performed to qualify the Si photonics interposer after processing. In the current design, a conservative KOZ of 100 μm was used for placing photonic components around the TSVs. As a result, the electrical performance of the TSVs can be tested after backside processing utilizing a multiple TSV chain test structure as well as a single TSV (Kelvin) structure at the wafer-scale, using a 4-point probe measurement. The TSV Kelvin and TSV multi-chain top-down and cross-section schematics, as well as typical measurement values of the respective resistance, are shown in Fig. 11.11, respectively [16]. TSV Kelvin resistance values of around 0.030 Ohm were obtained with good uniformity. TSV resistance of 0.105 Ohm/link was extracted from the chain test structure with 200 links, dominated primarily by the thin M1 contribution of ~0.07 Ohm/link.

Optical loss measured on a strip silicon waveguide and typical advanced grating coupler insertion loss measured on the silicon photonics interposer with TSVs are shown in Fig. 11.12. The propagation loss on wafers without and with substrate

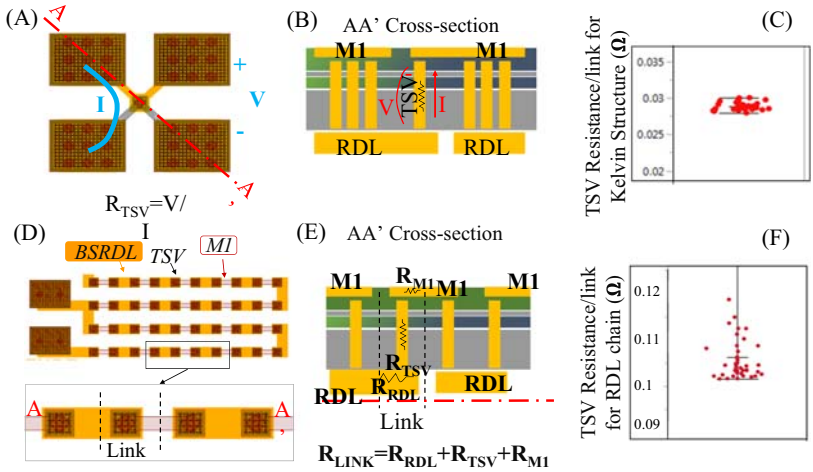

Figure 11.11 Schematic representation of M1-TSV-Backside redistribution layer Kelvin test structure: (A) top-down, (B) cross-section, and (C) extracted resistance per link for Kelvin test structure. Schematic views of M1-through silicon via-backside redistribution layer chain test structures: (D) top-down, (E) cross-section, and (F) extracted resistance per link for redistribution layer chain.

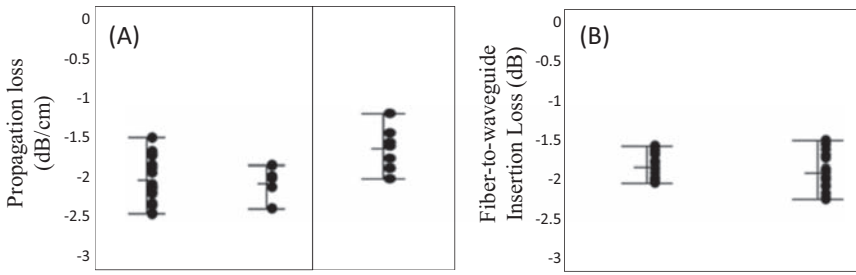

Figure 11.12 Measured response of the optical components on wafers with through silicon vias: (A) propagation loss before and after substrate thinning and (B) grating coupler insertion loss.

thinning was similar and within the wafer-to-wafer standard deviation of 0.3 dB/cm, varying from 1.9 dB/cm for the O-band on the wafer with full substrate thickness to 1.65 dB/cm after thinning. Typical spectra of a C-band ring modulator of 5-μm radius and an O-band ring modulator of 7.5 μm radius are shown in Fig. 11.13. The modulation efficiency and Q-factor of these ring modulators agreed with simulations, indicating no impact of the TSV processing on the structures as shown in Fig. 11.13. Fig. 11.14 shows the responsivity and dark current of typical Ge PD on a silicon photonics interposer of 100-μm substrate thickness. Again, no impact of the TSV processing on the photodetector performance is observed.

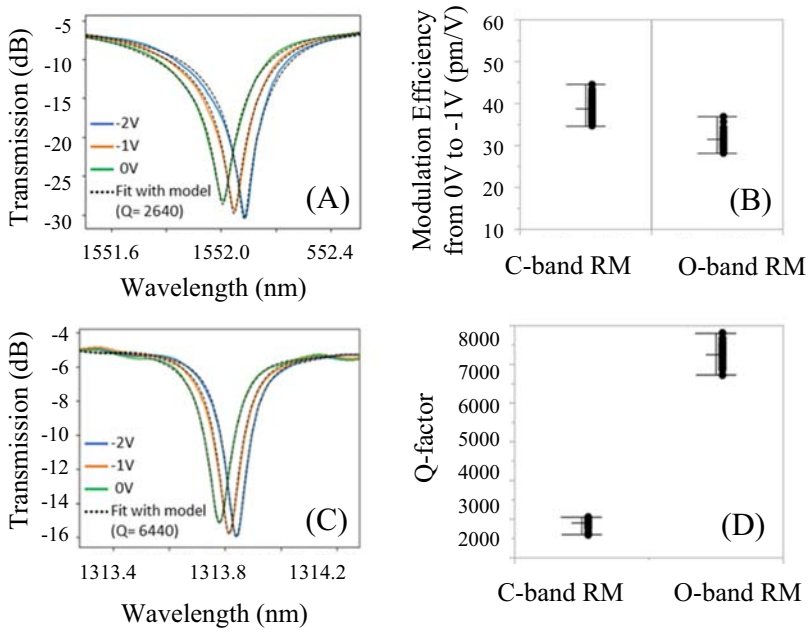

Figure 11.13 Ring modulator transmission spectrum: (A) at C-band and (C) at O-band. (B) Modulation efficiency and (D) Q-factor of microring modulator.

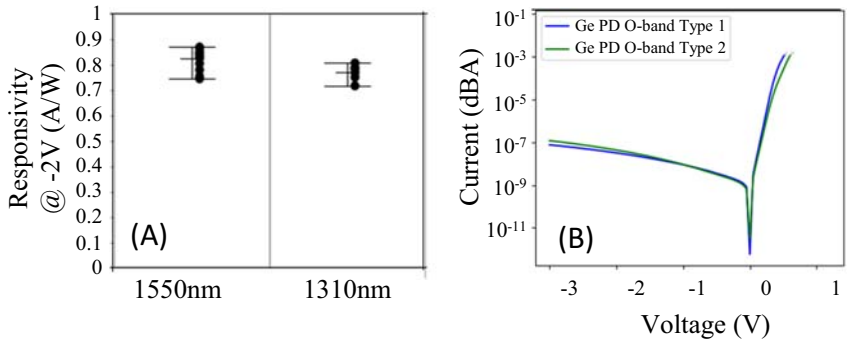

Figure 11.14 Measured characteristics of Ge photodetector on silicon (A) responsivity and (B) I−V characteristics.

11.5.1.3 Hybrid integration: key benefits and challenges

The development of the 10×100 μm TSV-middle module on photonics SOI wafers imposes some challenges, primarily related to TSV Cu filling and wafer warpage management throughout the integration flow [16]. First, the 2000-nm-thick BOX

layer combined with the ~1000-nm-thick top oxide cladding necessitates a thick resist during the dry etch of the TSVs. In addition, when transferring processes developed on bulk Si wafers to SOI wafers, some modifications in the processing steps are required compared to TSV-middle fabrication on standard Si substrates, for example, to deal with the increased stress owing to the dielectric SOI BOX layer stack thickness. Another challenge is void-free TSV Cu plating. A key parameter in this regard is the TSV density. Appropriate TSV densities to ensure good filling vary between 0.8% and 1.2%. Lower TSV density requires slowing down the plating process to achieve void-free TSVs, resulting in exceedingly long plating times.

As proof of principle, a TSV-enabled hybrid SiPh-CMOS transceiver prototype was designed, fabricated, and tested. The prototype employed 30 GHz Si ring modulators at the transmitter (TX) side and 50-GHz Ge PD at the receiver (RX) side, operating in the O-band and cointegrated with Si microring resonators-based WDM filters and TSVs in a 100-μm thick SiPh interposer, as shown in Fig. 11.15. In addition, Cu microbumps on a 50-μm pitch were implemented to connect the SiPh die to a codesigned 14-nm FinFET CMOS chip containing power-efficient modulator drivers and TIAs capable of operation at 40-Gb/s NRZ [21,22]. Owing to the tight microbump pitch and the compact footprint of the optical devices and electrical circuits, the TX and RX cells occupy each less than 100×300 μm^2, enabling a pathway to optical transceivers with Tb/s/mm^2 areal bandwidth density.

After assembly on a test PCB, a fiber array was actively aligned to the SiPh interposer. The assembly was subsequently powered and controlled through the TSV interconnects. At the same time, RF input and output signals could also be routed to the package through the TSV interconnects. However, for the reported results in Figs. 11.16 and 11.17, RF probes were used to deliver RF input data to

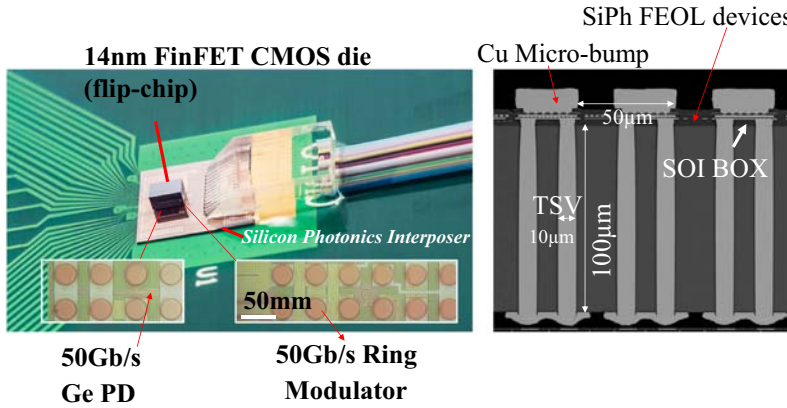

Figure 11.15 Hybrid SiPh-CMOS finFET transceiver prototype. Left: the prototype assembly with a CMOS die flip-chipped on the TSV-enabled SiPh interposer and attached fiber array. Insets show the Si-Mr modulator-based transmitter and Ge PD-based receiver unit cells, both connected to the CMOS chip through Cu-microbumps on a 50-μm pitch. Right: cross section SEM image of the 100-μm-thick SiPh interposer, showing through silicon vias and Cu microbumps cointegrated with SiPh devices on a 300-mm SOI wafer [21].

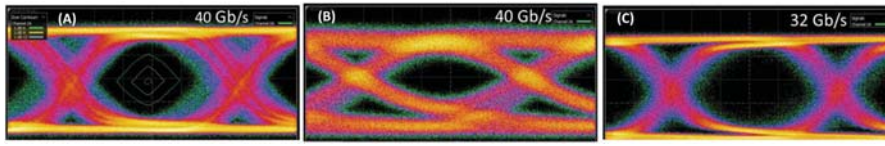

Figure 11.16 High-speed eye diagram generated by the hybrid CMOS-SiPh transceiver prototype: (A) 40-Gb/s transmitter with 4.6-dB extinction ration, (B) 40-Gb/s receiver, and (C) transmitter-to-receiver loop-back test at 32 Gb/s operating with 13-dBm laser input power and without optical amplification in the optical link [21].

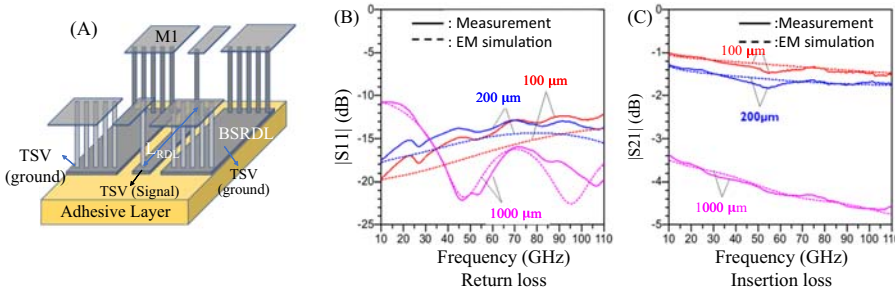

Figure 11.17 (A) Schematic representation of the through silicon via RF test structure. Measured and simulated RF response of the through silicon via test structure, as a function of back-side redistribution line length L_{RDL}, (B) return loss, and (C) insertion loss. RF insertion loss is below 1.5 dB at 110 GHz for LRDL = 100 µm. Solid lines for measurement data, and dotted lines for EM simulation data.

the transmitter side of the SiPh interposer and to probe the output of the TIA circuits. For the TX and RX cells, open eye diagrams were obtained at 40 Gb/s, as shown in Fig. 11.16. In addition, in a loop-back configuration without external optical amplification, an open-eye diagram was obtained at 32 Gb/s. Owing to the low parasitic capacitance of the microbumps and the optical devices, the modulator driver operated with 160 fJ/bit power efficiency and the TIA at 80 fJ/bit. In a fully functional link, heater control power and the required laser wall-plug power would raise the overall optical link power to 2−3 pJ/bit (excluding SerDes, CDR and PLL contributions).

Although in the presented demonstration, the TSVs were not used for high-speed data transmission, they feature remarkably high bandwidths of >100 GHz when implemented in a GSG transmission line configuration, opening a pathway for low-loss transmission of 100Gbaud data, as shown in Fig. 11.17 and as reported in [23]. Such large transmission bandwidth provides a pathway to electrical *I/O* lane rates up to 200 Gb/s.

11.5.2 Heterogeneous photonics integration platform

In a heterogeneous photonics integration scheme, the III−V lasers are processed simultaneously with the rest of the silicon photonics components. However, the

electronics drivers, amplifiers, and ASIC components are separately fabricated and packaged. Since the lasers are patterned along with the rest of the components, the light coupling loss is very minimal in this scheme. Heterogeneous fabrication of III lasers entails either bonding of the III−V chiplet or heteroepitaxial growth of the III−V materials onto the silicon. This scheme of integrating two different materials with different processes and thermal characteristics requires addressing several critical challenges in enabling an efficient laser (and semiconductor optical amplifier [SOA]) on a silicon platform. It is imperative to design an optimized III−V material stack, preferably using quantum materials, such as quantum dots, as the lasing source. These lasers integrated onto silicon waveguides could be either butt coupled or evanescently coupled. Due to better light coupling, power-efficient hybrid laser without the need for package assembly is a primary advantage of a heterogeneous integration scheme. This is not an easily implementable integration scheme, and Intel Corporation has done extensive research and adapted the technology for production. Thus the heterogeneous integration approach entails quasimonolithic integration of lasers and SOAs with silicon photonic devices (Fig. 11.18) [24]. The key advantage of this integration scheme is that it enables complete optical functionality, including III−V laser to be fabricated directly on silicon wafers, taking full advantage of the massive existing infrastructure of 300-mm silicon processing, wafer-level testing, and assembly, and scalability to very high-volume manufacturing. By combining all transmit and receive photonic functionalities onto a single wafer, fully integrated optical systems can be realized on a single die. In recent years this has become an established commercialized process [25].

11.5.2.1 Design considerations

Unlike other silicon photonics process technologies without III−V integration, the heterogeneous integration process must be designed to cooptimize the active III−V/Si laser and the passive and active silicon devices. One key distinguishing factor is in the choice of SOI wafers. While most other silicon photonics processes utilize a 220-nm silicon layer over a 2-μm BOX layer, the heterogeneous laser requires high

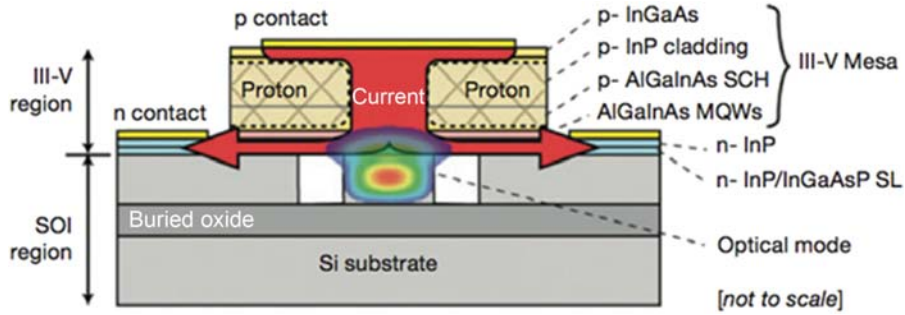

Figure 11.18 Cross-section of the laser light source used in the heterogeneous integration approach [26].

thermal conductivity for efficient heat dissipation. Therefore the BOX is reduced to 1um thickness. Typically, all active III−V devices, including lasers and SOAs, are designed in a region of each die where the gain medium is bonded, and all passive silicon devices are fabricated in the surrounding regions. In some cases, the silicon waveguide layer may be thicker in the laser regions (e.g., 400 nm) but thinner in the silicon passive/active regions to realize improved performance, such as faster phase modulation, with adiabatic waveguide couplers connecting between the devices in different regions. The heterogeneous laser device is typically on the order of several hundred microns long by a few tens of microns wide, so arrays of multiple lasers, SOAs, and III−V PDs can be designed with fine pitch alongside the other optical devices (Fig. 11.19).

11.5.2.2 Integration approach

The process for making the laser starts with the heterogeneous integration of an unpatterned stack of epitaxially grown quantum well material on a III−V wafer. Due to the mismatch of wafer sizes (300 mm silicon wafers versus 2−4″ InP wafers), a III−V die to Si wafer bonding approach is typically used. The generalized process flow is shown in Fig. 11.18. First, standard silicon photonics processing creates passive and active silicon components, including waveguides, silicon passives, modulators, and Ge PDs. Additionally, silicon nitride layers are typically fabricated for additional waveguide routing, spot-size converters, etc. Next, a rectangular region of a few mm on a side is reserved for the lasers to be fabricated. In this region, silicon waveguides and laser gratings are fabricated. Once all high-temperature steps are

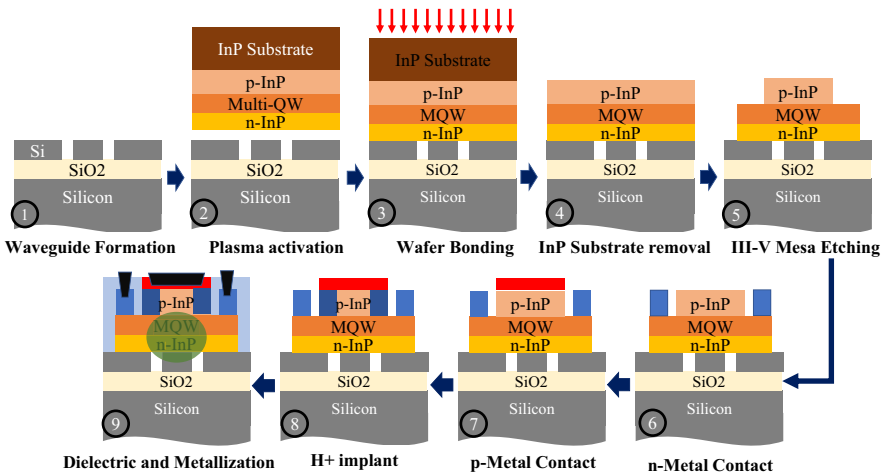

Figure 11.19 Heterogenous integration and fabrication of III−V/Si lasers in a standard 300-mm silicon fabrication line. (1) waveguide formation, (2) plasma activation, (3) wafer bonding, (4) InP substrate removal, (5) III−V mesa etching, (6) n-metal contact, (7) p-metal contact, (8) H + implant, and (9) dielectrics and metallization.

completed, the resulting dielectric stack is etched away, exposing the silicon waveguides and gratings in the laser region. A die-to-wafer approach is then used to bond a full unpatterned stack of III−V material onto the silicon in this region. This is typically done using plasma activation followed by a low temperature (<350°C) anneal. The InP substrate is removed mechanically and chemically, stopping on the p-layers of the epi stack. Once the unpatterned III−V stack is fully integrated onto the Si photonic wafer, the laser is fabricated [25]. First, the laser mesa and taper structures are made by patterning and etching the III−V material, followed by ohmic contact metallization to the p and n layers. Often, proton implantation is then applied to define the current channel, and dielectrics and metals are patterned to complete the electrical routing.

11.5.2.3 Hardware performance

The core of the heterogeneous laser integration approach consists of a III−V gain medium bonded directly to a silicon waveguide. This approach enables extremely low insertion loss (less than 0.5 dB, compared with >3 dB for discrete lasers [27]) because the laser is directly integrated into the silicon, with its cavity defined by gratings formed within the silicon waveguide, as shown in Fig. 11.20.

The heterogeneous laser integration approach has demonstrated good performance for many data center applications. Multiple wavelengths can be integrated on a single silicon chip by bonding different III−V materials and adjusting the gratings on the silicon wafer (see Fig. 11.21). Uniform performance has been demonstrated for 8 channels spanning wavelengths from 1270 to 1410 nm, offering the potential to multiplex 8 CWDM wavelengths in a single transceiver. Low voltage (<2 V) operation with 10 mW output power at 60 mA current and 80°C for 100 Gb applications and 25 mW output power at 100 mA current and 80°C for 400 Gb applications have been reported [27]. In addition, excellent long-term reliability is

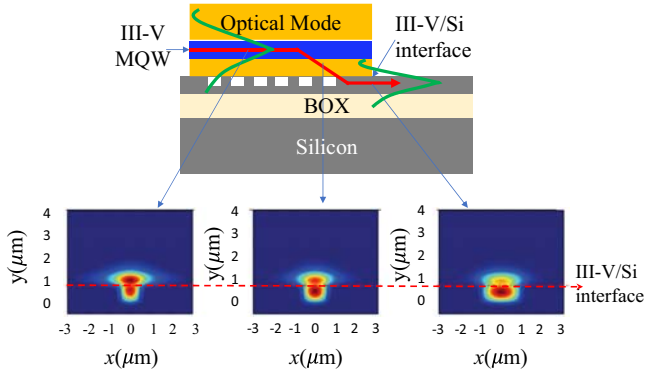

Figure 11.20 Schematic cross-section of a heterogeneously integrated laser on silicon-on-insulator platform, the gain medium is directly bonded to the device silicon, while the cavity is formed by gratings in the silicon waveguide [27]. Inset figures show the simulated mode profiles at various positions of the cavity.

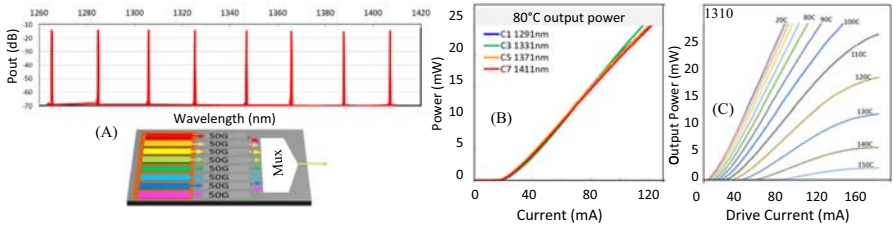

Figure 11.21 (A) Multiple wavelengths of hybrid lasers integrated on a single silicon die, (B) laser output power as function of drive current, and (C) laser output power.

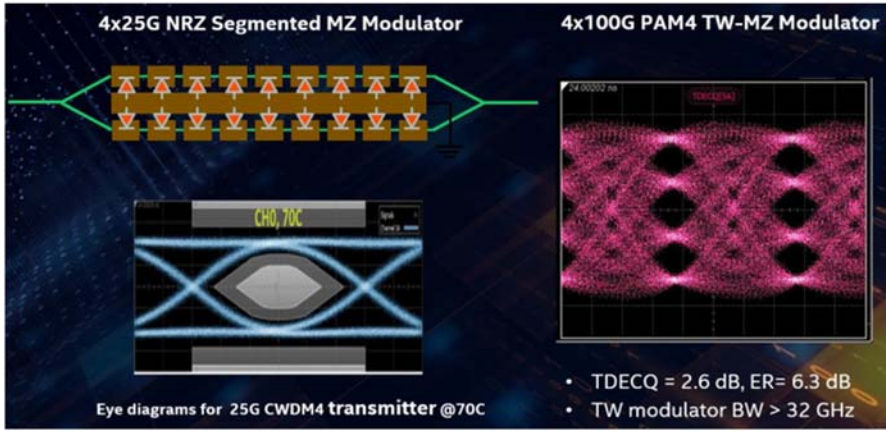
4x25G NRZ Segmented MZ Modulator
4x100G PAM4 TW-MZ Modulator
CH0, 70C
• TDECQ = 2.6 dB, ER= 6.3 dB
• TW modulator BW > 32 GHz
Eye diagrams for 25G CWDM4 transmitter @70C

Figure 11.22 Performance of two types of heterogeneous integrated silicon modulators: 4×25G NRZ MZM used in Intel's 100-G CWDM4 products, and 4×100G PAM4 TW-MZM used in Intel's 400G Dr4 product [25].

achieved due to the absence of laser facets, smaller junction temperature increases, and $\sim$5X lower current densities compared to other typical III−V based directly modulated lasers (DMLs) [28].

The integrated laser is particularly compelling in combination with silicon modulator technologies, which can be made using standard silicon photonics processes. Several types of modulators have been demonstrated in the heterogeneous integration platform, and many are shipping in high volume. These include 4×25G NRZ segmented Mach−Zehnder modulators (MZMs) and 4×100G PAM4 traveling wave (TW) MZMs, as shown in Fig. 11.22.

11.5.2.4 Heterogeneous integration: key benefits and challenges

Being an indirect bandgap semiconductor, silicon is notorious for not being able to generate light. However, it is highly capable of optical functions, including low loss waveguides, high speed modulation, and efficient detection using epitaxially grown Ge PD. Heterogeneous integration thus offers several distinct advantages,

the most important of which is the ability to realize a complete photonic system on a chip. The laser integration enables the perfect marriage of the direct bandgap III−V materials to enable optical gain with the exceptional properties of silicon as an optical material in the infrared. Due to practical process incompatibilities, III−V laser integration makes it challenging to coprocess the photonics with a fully standard CMOS process. However, one may argue that the heterogenous integration approach of separating the photonics from the circuitry allows both to be individually optimized (i.e., the circuitry can leverage the most advanced CMOS processes) to get the best of both worlds. At the same time, the complexity and hence yield and cost of the process are improved by not monolithically integrating the photonics with the electronics. Finally, it can be advantageous to select the best stand-alone electronic chips from established processes with very fast cycle times available in the industry.

However, there are also several unique challenges associated with the heterogeneous integration approach: For instance, laser fabrication creates significant topography, which makes typical planarized Damascene interconnect quite challenging. Therefore lithographic patterning of metal lines is more attractive despite its known drawbacks associated with topography. Also, the addition of an active device late in the process flow places additional requirements on metallization since contamination with Cu or Au is likely to cause reliability problems; therefore additional barrier layers may be required to prevent metal diffusion into the active regions. Finally, the stress of the silicon wafer and die must be carefully managed throughout the fabrication and assembly processes to ensure no adverse stress-related impacts on the laser.

11.5.3 Monolithic photonics integration platform

A monolithic integration platform enables the cointegration of electronics' driver circuits and photonic devices on the same silicon substrate using mature CMOS process flow. The most significant advantage of this close integration is a significant reduction of parasitics that exists in the conventional hybrid approach, which results in improved energy efficiency [29]. Monolithic integration not only gives this performance boost but also a commercial boost by removing the need for advanced packaging of electronics and photonics and the associated cost.

Even though the electronics drivers and amplifiers are integrated into the silicon platform, the lasers and ASICs do need to be packaged separately at the wafer level or through an interposer. Unlike other photonic integration solutions, such as hybrid integration, monolithic SiPh technology requires process optimizations for photonics and electronics to perform simultaneously, leading to additional device performance compromises. IBM conducted a substantial amount of research on monolithic silicon photonics, which GlobalFoundries (GF) later adapted for its 90-nm technology. GF has also developed a 45-nm RF-CMOS silicon photonics technology that is competitive in performance. The Leibniz Institute for High-Performance Microelectronics (IHP) also offers low-volume manufacturing of monolithic SiPh technology in its BiCMOS platform. In a nutshell, a monolithic approach in which control electronics

and photonic transceivers are integrated in the same process takes the process technology to the next level and can provide true parasitic-less integration, thereby enabling the fabrication of faster and more efficient transceivers and providing significant performance and cost-saving benefits.

11.5.3.1 Design considerations

In this section, we will discuss some of the primary parameters, such as (1) thermal isolation, (2) dual silicon thickness, and (3) backend metal thickness, that must be considered for designing a monolithically integrated circuit.

Monolithically integrated optical and CMOS devices can more directly influence each other thermally than in the hybrid platform due to their significantly closer proximity. Thermally tuned optical devices, such as microring modulators or ring-based WDM filters, locally generate a gradient of temperature that may thermally affect neighboring devices. In addition, in a monolithic platform, any CMOS circuitry, especially a high-speed logic or a driver or a TIA, is a source of heat that can thermally affect susceptible silicon photonic devices. The heat generated by monolithic CMOS circuitry is typically higher than BULK or even standard SOI CMOS circuitry due to the significantly thicker BOX needed for photonics (typically around 2 μm), which impedes thermal conduction through the thick silicon substrate. Thus thermal budget and thermal cross talk in the monolithic platform are serious concerns not only from a system-level point of view but also in terms of quality and reliability. A CMOS designer may expect a recommended maximum current density that can flow through a transistor or even gate pitch spacing for specific types of design from the PDK. From the system level point of view, thermal crosstalk between any two devices may negatively affect the final performance. Hence a minimal separation or distance between a thermally sensitive silicon device and potential thermal aggressor (heater or CMOS circuit) must be considered in the design. Another design consideration offered by many advanced photonic platforms is to locally remove the silicon substrate underneath the BOX to isolate devices thermally. Such a fabrication step is helpful in thermally tuned silicon devices to improve thermal tuning efficiency. By removing the silicon around the driver electronics circuitry, the thermal resistance can be increased. This approach can also be used to isolate optical devices thermally [30]. In addition, CMOS compatible materials with negative thermal coefficients can be integrated to isolate devices. On top of that, a chip designer can take advantage of metal stack with very low thermal resistance to build a heat sink that will further thermally isolate. Thus there is a potential risk of overheating both electronics and photonics devices in a monolithic integration scheme. However, this (overheating of the transistors or the neighboring photonic devices) can be effectively mitigated by following the design for manufacturing guidance from a PDK provided by the foundry (e.g., GF 45SPCLO PDK, silicon photonics PDK on a 45 nm SOI technology node and caters to C, L, and O bands).

It is essential to control the thickness of silicon in both the photonics and electronics regions. The incoming wafer surface and thickness uniformity are imperative to maintain. During the process, oxidation and etching contribute to variability.

Therefore it is vital to start with thick silicon and reduce it to an optimum thickness. In monolithic photonics, the incoming silicon is sub-200 nm, and the final photonics region is even thinner [29,31]. Compared to hybrid silicon photonics, a thin crystalline Si layer adopted in the monolithic SiPh technology presents a few design challenges for enabling functional photonic devices. One potential limitation is the weakened field confinement for components working in transverse magnetic (TM) modes, especially when the operating wavelength shifts towards C and L bands [13]. The resultant mode leakage loss may lead to a noticeable deterioration in device performance. A feasible solution to overcome such a limitation is to utilize both the Si and SiN layers on the monolithic platform to create a stacked waveguide for tightly confined TM modes [13]. Without introducing additional process customizations, the optical leakage loss can be significantly reduced using the stacked Si−SiN designs, which opens new opportunities to enable a powerful class of high-performance active and passive optical devices in a foundry platform.

An ideal backend-of-line in an electronics circuit requires only a few metal layers to reduce the parasitics. However, from the dense logic point of view, several metal layers are required to ensure the successful synthesis of functional logic. In addition, thick metal layers are needed for proper power distribution. The thicker metal layers also provide better RF capabilities to the system. For example, passive electronics devices, such as high Q-factor inductors or transmission lines, can be realized on thick metal layers to ensure high-quality performance and low signal insertion loss [31].

In general, the design procedure is remarkably similar in both a hybrid and a monolithic CMOS-SiPh foundry. It is important to keep in mind the above-mentioned key differences (CMOS-SiPh proximity, thermal crosstalk etc.) to design reliable electro-optical circuits.

11.5.3.2 Integration approaches

High volume manufacturing of the monolithic SiPh integration approach employed by GF leverages a state-of-the-art 45 nm RF-SOI CMOS technology [29,31]. It enables optical product designs with comprehensive C, L, and O band photonic device libraries with an option to monolithically integrate active and passive circuitry for analog and digital control. A typical process flow and schematic cross-section of the CMOS silicon photonics platform are shown in Figs. 11.23 and 11.24, respectively. It is fabricated on a 300-mm SOI wafer with 2 um thick buffer oxide. The SOI silicon has a dual thickness, 160-nm-thick silicon for photonic components

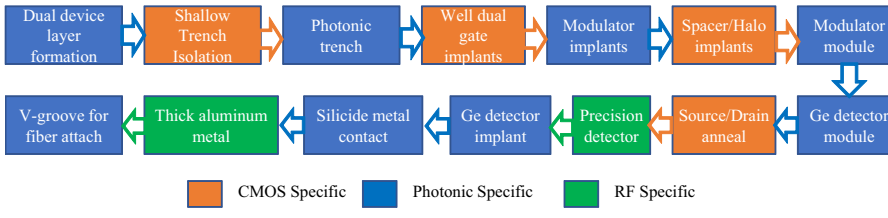

Figure 11.23 Critical process module for the monolithic silicon photonic technology.

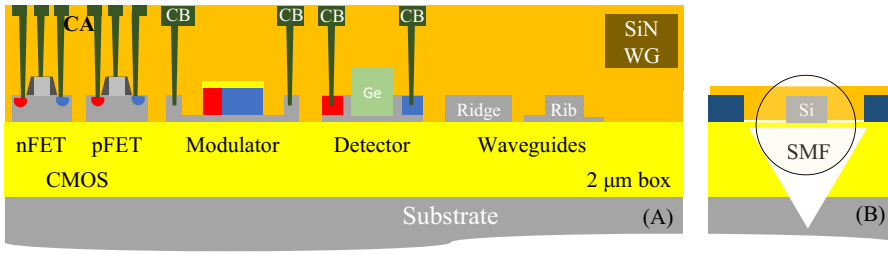

Figure 11.24 Schematic cross-section of (A) monolithic silicon photonics technology with CMOS electronics and photonic devices and (B) V-groove for input/output fiber attachment.

and 88-nm thin silicon for electronics. The technology provides high-quality active and passive silicon photonics devices with high-fT CMOS transistors by employing dual silicon thickness and two contact modules CA and CB. The 45-nm manufacturing process leverages 300-mm immersion lithography and OPC for patterning waveguides to optimize loss and device performance across various device structures. Compared to a dry lithography patterning process, immersion lithography provides higher performance and lower optical loss [32]. The 45-nm photonics technology node thus anticipates supporting high-performing optical transceivers pushing modulation speed up to 100 Gbd [31].

The passive device offerings incorporate fully etched ridge silicon waveguides, partially etched rib silicon waveguides, and multimode silicon waveguides. In addition, SiN-based passive devices including waveguides are fabricated on the same platform [33]. The waveguide loss in a monolithic SiPh technology is comparable to hybrid SiPh solutions, allowing full realization of the system benefit of monolithic integration of photonic and electronic devices. The active photonic devices are (1) high bandwidth modulators, such as MZMs and ring modulators, and (2) epitaxial germanium photodiode. In addition, the technology supports (1) coarse wavelength division multiplexing (CWDM), (2) high-efficiency isolated thermal phase shifters, (3) RF passives, (4) ESD protection, (5) E-fuse embedded memory, and (6) standard cell digital library. Silicon is an indirect bandgap material. Consequently, light-source (laser) must be integrated using alternative means (hybrid or heterogeneous through bonding or monolithic through epitaxial integration). On-chip lasers are typically constructed out of III−V materials. As these materials are not CMOS compatible, they are challenging to integrate in the monolithic process line. In addition, the optical interface (input/output) is of paramount concern as the grating-based couplers are highly lossy and wavelength dependent. Edge couplers using subwavelength grating structures are found to be a potential solution to solve optical coupling issues with the photonics chips. Subsequent sections briefly discuss the integration aspects of *I/O* couplers and lasers.

11.5.3.3 I/O edge coupler integration—V-groove

A monolithic SiPh platform incorporating V-groove technology requires several significant changes to the best-known methods. As CMOS technology moved to

smaller nodes, a focus was placed on improving patterning capabilities and enabling small patterns on flat wafers. Photonics and the addition of V-grooves change this dynamic and require processes that can pattern images over large topography. These V-groove structures are large, and are preferentially etched along the [111] plane of the Si [100] substrate and are formed on the edge of the silicon die. With thick resists and deep cavities, the etching process poses significant challenges to alignment systems and uniformity control. It has been identified that the tetramethylammonium hydroxide etch is one of the most effective methods for addressing these challenges and creating a deep, crystallographic etch. Etching is halted when two Si [111] planes come into contact with one another. This is due to the fact that the Si [111] plane etch rate is the slowest among all the other crystal planes, and the process step only exposes the Si [111] plane to the solution. The depth of the etch is determined by the size of the window that has been opened. The optical fiber center is then passively aligned to the spot-size converter that is fabricated on the PIC by inserting the fiber into the V-groove structure (Fig. 11.24B). Once the fiber is aligned, either epoxy or a index-matching fluid is dispensed onto the V-groove for improved mode matching and better mechanical stability of the fiber inside the groove. This provides a low-loss, low-cost *I/O* fiber attachment [34].

11.5.3.4 Laser integration

Monolithic (epitaxially grown and processed) or heterogeneous (die bonding and processing) integration of laser in a photonics−electronics sequential process flow is difficult within the current fab and integration constraints. Therefore just as in hybrid technology, laser components must be independently fabricated and packaged. The challenges associated with light source integration are thus associated with the coupling efficiency and relative intensity noise (RIN), or mode hopping. The hybrid laser assembly requires the creation of the cavity in conjunction with the V-groove process to enable fiber attached based optical *I/O*. The proper laser placement and alignment with photonic waveguides depend strongly on the precisely controlled cavity dimensions [35]. Additionally, the metal wiring layers, and solder connection must be fabricated within the laser cavity, providing mechanical support and electrical connections. The metallization wiring will often help with thermal stabilization, acting as a thermal radiator. The solder thickness must be carefully controlled to achieve consistent solder compression and good electrical contact during the laser attach and vertical alignment. A typical laser integration process tolerates no more than 2−3 μm of horizontal or vertical misalignment to ensure better than 3-dB coupling efficiency.

Once the laser is appropriately aligned with silicon photonics optical layers (silicon or silicon nitride) and all electrical connections to the laser diode are made, the coupling efficiency depends on the spot size converter structure. This part of the waveguide is responsible for converting the optical mode from a laser to the silicon or silicon nitride waveguide with minimal or no loss. Therefore comprehensive numerical modeling is needed to optimize the design of the silicon spot size converter.

11.5.3.5 Hardware performance

Table 11.1 compares the CMOS performance of the GF 45RFSOI and 45SPCLO technologies. The 45SPCLO technology offers significantly improved CMOS devices with an RF-friendly, 8-metal layer stack back-end. The first five copper metal layers are suitable for complex logic circuits. The top two metal layers include a 1.2-μm-thick copper and 1.2-μm aluminum for routing low loss transmission lines (Fig. 11.25A) and building high-Q inductors. The 45SPCLO platform also offers an advanced IOSMF coupling solution; the self-aligned V-groove and spot size converter provide a lower loss solution for fiber or fiber array attachment, providing robust, high-performance, and high-yield optical coupling.

The primary advantage of the 45SPCLO platform is to offer 45RF CMOS performance while still providing state-of-the-art photonic devices. To accomplish this, significant process optimization needs to be done to mitigate performance impact on the CMOS devices and some of the active photonic devices, such as germanium (Ge) photodiode. Ge is epitaxially deposited using optimized preclean conditions in

Table 11.1 45SPCLO CMOS technology comparison to the parent RF reference node 45RFSOI.

Attributes	45RFSOI	45CLO SipPh
Mask count	40	90
Tsi SOI/BOX thickness (nm)	88/200	88/2000 CMOS 160/2000 Photonics
Contacted gate pitch (nm)	190	250
Nominal supply (V)	1.0/1.8	1.0/1.8
Minimum physical gate length (nm)	41	41
Monolithic CMOS	RVT, UVT, AVT, DG, (1.8V)	RVT, UVT, AVT, DG, (1.8V)
NFET peak f_T/f_{MAX} (GHz)	300/355[a]	280/340[a]
PFET peak f_T/f_{MAX} (GHz)	245/305	230/290[a] with eSiGe

[a]SiPh integration impacts.

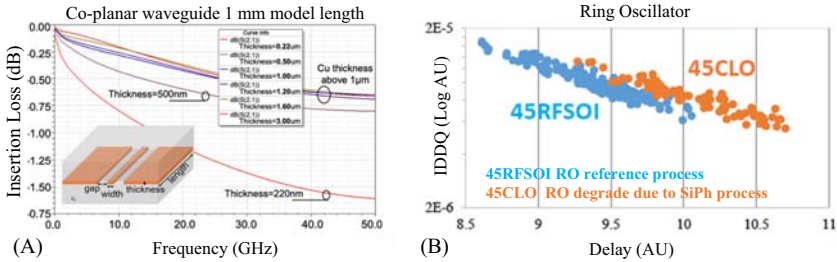

Figure 11.25 (A) Simulated insertion loss of a 1 mm long coplanar GSG transmission line for different conductor thickness. (B) Measured insertion distribution of the OFF-state leakage current IDDQ as a function of ring oscillator delay.

a shallow cavity that is patterned on the silicon. Ge integration imposes significant thermal constraints that includes maintaining the temperature below Ge melting temperature. This means that its deposition and implants must be done only after the CMOS source/drain anneal (980°C). The challenge now is that the CMOS dopants, such as boron, impose thermal constraints on the Ge process. Thus low-temperature Ge epi formation (780°C) is necessary to maintain the CMOS devices' performance while ensuring that single crystalline Ge is available for the photodetector formation. With this, Fig. 11.25B highlights that the CMOS performance is within 7% of the target device while keeping the high-performance characteristics of the photodiode.

As monolithic integration does not contribute parasitic load; the driver power efficiency, and the total modulation speed can be significantly improved. Even more substantial impact on the system performance can be observed on the optical receiver, where the 3-dB bandwidth of the TIA strongly depends on the total input capacitance. Also, the total input load can be further reduced by removing the ESD protection because, in monolithic integration schemes, the transistor gates do not need protection. Therefore the total input capacitance of the monolithically integrated receiver could be in the range of 20fF, including photodiode, TIA input stage, and metal interconnect. Even though the transistor speed in the 45SPCLO technology, measured by the fT parameter, is comparable to the 14-nm FinFET, the achievable modulation speed of 45SPCLO is expected to be twice that of 14-nm FinFET for optical transceiver applications. Owing to the monolithic integration of silicon photonics and CMOS devices as well as the RF friendly back end (high fmax), the superiority of the 45SPCLO technology for short-reach optical interconnects in comparison to the hybrid approach is summarized in Table 11.2 [31].

The performance figures of the photonic building blocks are measured to be very promising. Fully etched single-mode silicon ridge waveguide loss is 1.032 dB/cm. Whereas the propagation loss of the multimode silicon ridge waveguide is measured to be 0.632 dB/cm. Silicon nitride waveguide loss is 0.39 dB/cm. Bend loss for a 90-degree Si bend with 10 μm radius of curvature is nearly 0.005 dB. The insertion loss

Table 11.2 CMOS technology comparison for low-power silicon photonics transceiver applications.

Parameters	**14nmFF**	**22FDX**	**BiCMOS**	**45CLO**
NMOS *f* (GHz)	270	350	500	280
Supply voltage (V)	0.8	0.8	3.3	1.0−1.1
Technology	BULK	SOI	SiGe HBT	SOI
Mx res and cap	High	Low	Very low	Very low
Integration	Hybrid	Hybrid	Hybrid	Monolithic
Parasitic load (fF)	5−30	5−30	5−30	2−3
ESD diode (fF)	∼50	∼50	∼50	0
Estimated TIA 3 dB BW@1 kΩ gain	24 GHz* (70 fF)	28 GHz* (70 fF)	33 GHz* (70 fF)	47 GHz* (20 fF)

Source: Adapted and modified from R. Herrick, Reliability qualification and failure mechanisms for semiconductor lasers and fiber optic transceivers, in: Optical Fiber Communication Conference: Optica Publishing Group, 2019 [28].

of a wavelength independent coupler with 50:50 splitting ratio is ∼ 0.07. This technology also provides polarization splitter/rotator (PSR) with a hybrid waveguide configuration of silicon and silicon nitride. The insertion loss for TM and TE PSR are measured to be 0.41 and 0.65 dB, respectively. High speed modulator with MZI configuration offers a 3-dB bandwidth up to 35 GHz with −2-V bias voltage. Ge photodiodes with 0.7-μm epi layer gives a 3 dB bandwidth of 47 GHz, responsivity of 0.925 A/W with a dark current below 28 nA. 45SPCLO technology also supports input/output fiber attachment (V-groove) with a maximum insertion loss of −0.84 and −2.1 dB for TE and TM polarizations, respectively. The optical return loss for the inline-fiber attachment is measured well below −35 dB. Table 11.3 summarizes the performance figures of the photonics building blocks in GF 45SPCLO PDK. As improvements to the process flow happen, and as PDK is updated, these numbers will also get revised.

11.5.3.6 Monolithic integration: key benefits and challenges

Monolithic silicon photonics technology promises a very energy-efficient optical link that can successfully compete and replace conventional copper-based interconnects. The CMOS control/driving circuitry must be codesigned with photonic devices to maximize power utilization and link efficiency. On the transmitter side, the electronic driver must provide the required voltage swing to the electro-optical modulator. The power consumption of such a driver is proportional to the modulation speed. It depends on the total capacitive or resistive load that must be reloaded to encode every single bit of information. In the advanced application based on microring modulator or germanium-based electro-absorption modulators, the monolithic SiPh platform has a considerable advantage. For example, a typical flip-chip

Table 11.3 Measured response of the important photonics components fabricated using GF45SPCLO technology.

Parameter	Measured response
Silicon ridge waveguide loss	1.032 dB/cm
Silicon rib waveguide loss	0.725 dB/cm
Multimode silicon ridge waveguide loss	0.632 dB/cm
Silicon nitride ridge waveguide loss	0.39 dB/cm
Multimode silicon nitride ridge waveguide loss	0.299 dB/cm
Silicon bend (radius = 10 μm) loss	0.005 dB/bend
Wavelength independent 50−50 coupler loss	0.07 dB/device
Polarization splitter rotator TE/TM output loss	0.41/0.565 dB/device
Mach−Zehnder modulator 3-dB bandwidth (−2 V)	35 GHz
Germanium photodiode responsivity	0.925 A/W
Germanium photodiode 3-dB bandwidth	47 GHz
Germanium photodiode dark current	28 nA

Source: Adapted and modified from R. Herrick, Reliability qualification and failure mechanisms for semiconductor lasers and fiber optic transceivers, in: Optical Fiber Communication Conference: Optica Publishing Group, 2019.

integration parasitic is comparable to the device load (~50 fF), so the energy-saving of monolithic integration that removes parasitics can be significant.

Furthermore, the monolithic process removes unwanted integration load and parasitic capacitance associated with ESD protection on the receiver side. By reducing total input capacitance on the TIA, the designer can achieve a higher gain with the same bandwidth or higher bandwidth with the same gain. Codesigning CMOS with photonics devices also means proper optimization of CMOS circuitry to meet application specifications.

In most cases, the electrical bandwidth on the TIA should match the optoelectrical bandwidth of the detector to minimize noise amplification; however, there might be some exceptions from this rule. For example, the high data-rate application typically requires high modulation speeds hence high or extremely high bandwidth drivers and receiver circuitry. In addition, to overcome optical device speed limitation, the amplification block of the receiver may have bandwidth extension circuitry to recover the data at the cost of excess noise.

The ability to cosimulate photonics and electronic devices simultaneously is one of the essential features provided in the PDK. The photonic devices have a fundamentally different operating nature than transistors, and the standard simulation tool for CMOS cannot model photonic devices in the optical domain. However, some shortcuts can be implemented to enable cosimulation, hence codesign. One of the most common approaches is to model electro-optical behavior in a higher-level language used during CMOS simulations. Here Verilog-A is widely used to model both the electrical and optical nature of photonic devices. Such a Verilog-A model can approximate the optical response, but the accuracy may be outstanding only when the model parameter range is limited. The active photonic devices, such as modulators or PDs, can be accurately and efficiently cosimulated with corresponding CMOS circuitry providing reliable feedback data to the system designers.

On the other hand, passive devices usually have straightforward models associated with their loss. The biggest challenge in cosimulating complex systems is associated with different light polarization or simultaneous simulations of more than one operating wavelength. Those additional dimensions in the optical domain must be captured in the device models for accuracy.

11.5.4 InP III–V photonics integration platform

With their direct bandgap, InP III–V materials provide an advantage to the PIC by integrating an epitaxially grown laser source with other photonic components. While the laser source integration can be thought of as monolithic, the electronics must be separately fabricated and packaged, thus leading to a copackaged hybrid integration scheme.

InP is an ideal material choice for photonic integration. InP's ternary and quaternary compounds exhibit several attractive properties, including gain, transparency, absorption, detection, and electro-optical modulation. In addition, strain engineering on quantum well structures enable light emission spanning wavelengths from 0.9 to 2 μm, covering the frequency spectrum relevant for optical communication (1.26–1.625 μm).

Also, InP has high electron mobility, making it preferred for high-speed electronic devices. However, besides special-purpose electronic devices, such as high electron mobility transistors (HEMTs) and heterojunction bipolar transistors (HBTs), InP is not preferred for electronic circuits due to high manufacturing costs. For transceivers, the InP PIC platform still relies on silicon-germanium electronics for high-speed modulator drivers and TIAs. One potential downside of the InP process is its multilevel epitaxial steps for growth and regrowth, which generates high defect density in this platform. The following sections discuss the aspects of InP-based device design, fabrication, device performance, associated technology benefits, and challenges.

11.5.4.1 Design considerations

As a binary III−V semiconductor, InP has a crystallographic structure of two overlapping face-centered cubic lattices with a lattice constant of 5.87 Å. Due to its excellent electron mobility [36], InP is popular for high-speed applications. InP has a direct bandgap of 1.344 eV at room temperature; this corresponds to emission in the near infrared (932 nm). A major advantage of InP is its ability to serve as a substrate for epitaxial growth of other III−V semiconductors. As an example, ternary (InGaAs, InAlAs) and quaternary compounds (InGaAsP, AlInGaAs) can be created by combining two or more binary compounds (InP, GaAs, InAs) [37]. Although in bulk form, these materials are limited to the emission wavelengths corresponding to their bandgap energies, in reduced dimensionality form (quantum wells/ quantum dots), the emission and absorption wavelengths can be tailored to any application in the optical communication bands. Adding strain to quantum wells allows the material system to cover any wavelength in the 0.9- to 2.0-μm window. An additional benefit is that the bandgap energy has an inverse relationship with the refractive index, allowing lower-index materials, such as InP, as cladding layers for higher index cores, such as InGaAsP or AlInGaAs. InP and associated ternary and quaternary compounds can achieve all required functional sections of a transmitter/ receiver PIC: light generation, amplification, modulation, detection, and passive optical routing.

InP lasers are typically based on narrow compressively strained multiquantum well (MQW) active regions with an effective bandgap detuning of ∼25 nm relative to the target operating wavelength. Examples of laser types include (1) distributed feedback (DFB) and (2) distributed Bragg reflector (DBR) variants. A widely used tunable DBR laser is the Sampled Grating DBR (SGDBR). This type of laser requires a passive/active integration scheme that consists of an active gain section, passive front and rear SGDBR mirrors, and a passive phase tuning section. The differentially tuned grating mirrors consist of reflection combs of different spectral spacing to achieve continuous tuning over the entire extended C-band (tuning range is typically over 45 nm). The lasers demonstrate a high side-mode suppression ratio (SMSR) across this tuning range (> 45 dB) with less than 2 dB power nonuniformity across the entire range without any active power flattening.

The optical modulator designs on InP use quantum wells with much larger detuning than the gain sections, allowing compact and power-efficient modulators

with lower loss, excellent linearity, high signal-to-noise ratios, and good reach. These devices are very efficient with typical V_π values in the order of $1.5 - 2.5$ V. Low $V\pi$ values imply that SiGe drivers can easily drive InP modulators. Also, the low V_π capability of InP with high electrical/RF bandwidths (>50 GHz) enables ultrahigh baud rates (>90 Gbaud). The photodetector building block can readily be implemented in InGaAs lattice matched to InP since it is a direct bandgap semiconductor. Passive functions, such as splitters, combiners, phase shifters, polarization splitters/rotators/combiners, 90-degree hybrids, waveguides, muxes, and demuxes can also be implemented in InP, with their performance competitive to those of other PIC platforms.

11.5.4.2 InP fabrication approach

InP PICs are fabricated using conventional InP processing techniques. A block diagram of Infinera's fabrication flow is shown in Fig. 11.26. Fabrication of the PIC begins with the formation of the multiple epitaxial layer structures required to integrate numerous active and passive devices. Epitaxial layers are grown using metal−organic vapor-phase epitaxy in multiwafer reactors. Active elements consist of MQW regions and passive sections consist of bulk heterostructure waveguides.

Several techniques have been developed over the years to create device structures with variable bandgaps within the same device. The most popular methods are selective area growth [38], etch-and-regrowth (butt-joint approach) and impurity-induced layer disordering (IILD). In selective area growth, the substrate is partially covered with an amorphous mask, such as SiO_2. The masking material redistributes the group III precursor elements in the lateral direction, leading to an in-plane

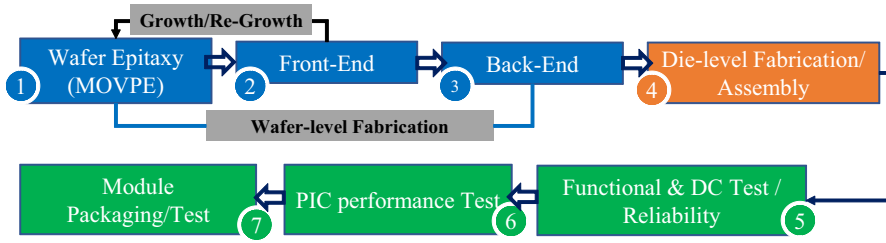

Figure 11.26 Schematic flow diagram of the sequence of primary manufacturing steps to fabricate and test commercial III−V PICs. (1), (2), and (3) are Wafer-level and (4) is die-level fabrication. (1 and 2) Active and passive component integration: distributed feedback laser, local oscillator, semiconductor optical amplifier, electroabsorption modulator, RF, pin photodiode, variable optical attenuator, arrayed waveguide, RF, and DC component growth/etch/regrowth integration. (3) Waveguide definition, interdevice and channel isolation, P/N contact formation, device passivation, DC and RF band pad Dep. (4) Singulation, AR-coat, and chip-on-carrier solder die-attach. (5) Functional and DC test: performance and wavelength stability screening. (6) Full PIC performance test—10 channel DC and RF test and link (BER) testing. (7) Full module performance test—module packaging and testing.

thickness modulation in the final structure. An advantage of this scheme with more than two group III precursors is that a single epitaxial process can modulate the composition, thickness, and bandgap by varying their growth rate factors. A disadvantage is that interfaces tend to be nonabrupt and passive structures will be lossy due to the presence of free carriers [39]. The second method relies on selective etching and subsequent regrowth of an optical device adjacent to an existing optical device structure [38,40]. Although this approach requires a patterned regrowth, the butt-joint process does provide maximum flexibility in choosing the properties of the two sections as the properties of one can be completely decoupled from the properties of the other, aside from the fact that they need to be nominally lattice-matched to InP. This flexibility adds fabrication complexity that needs to be managed to mitigate the impact on yield. A third technique to selectively change bandgaps is IILD [41]. This technique utilizes diffusion of impurities/point defects into the lattice to selectively alter the bandgap across the wafer. When applied to quantum wells, the technique is referred to as quantum well intermixing (QWI). During the process, a selective bandgap can be increased by a patterned implant and subsequent annealing steps to intermix the quantum wells and barriers in regions that are to become passive or active components. Bandgap shifts saturate at roughly 100 nm in InGaAsP, but intermediate shifts are harder to achieve and need to be carefully calibrated. Intermixing must also be done in the absence of other rapidly diffusing species, such as Zn, which means that QWI needs to be complete before growing any p-type cladding layers [42].

The epitaxy process requires control of composition, strain, doping, and thickness uniformity (within wafer, wafer to wafer, and lot to lot) similar to or exceeding that needed for the growth of vertical-cavity surface-emitting lasers. After epitaxial growth/regrowth and front-end wafer fabrication (patterning and etching) are complete, wafers are subjected to a back-end fabrication sequence similar to the processes used to form an HBT. The BEOL defines active/passive waveguides, interdevice and channel-to-channel electrical and thermal isolation using ion implantation and deep trenches, contact p and n regions of the active devices, Ohmic contact areas, DC and RF bond pad definition, and material passivation.

Dry etching, preferred for best pattern transfer and lower dependence on varying etch rates on different crystallographic planes and material compositions, is implemented using conventional reactive ion etch and high-density inductively coupled plasma tools. These tools allow the definition of structures that require precise dimension control, etch depth control, sidewall profile, and roughness, all requirements for precise wavelength control and low loss. Finally, conventional metallization and dielectric deposition techniques utilized to fabricate III−V optoelectronic devices are employed in the fabrication process.

After fabrication in wafer form, the wafers are subjected to die fabrication. Wafers are singulated, and individual dies are coated with an antireflection (AR) coating. The dies are subsequently die attached to a submount. Next, the resultant chip on carriers are subjected to a test and reliability screening/burn-in sequence to screen for performance and wavelength stability. The fabrication sequence is completed by a rigorous final test sequence that includes DC and RF testing.

Prior to the assembly of the PIC transceiver, the PIC chip on carriers are soldered on a thermoelectric cooler (TEC) and enclosed in a metal package. The PICs are fiber-coupled or free-space optics coupled, and a monolithic modulator driver and TIA array is electronically coupled to the devices. More details of the subassembly/packaging process can be found in Chapter 12. Many of the fabrication specifications—from epitaxy to etching—are highly demanding. However, through rigorous manufacturing process controls, process optimization, and continuous advancement of fabrication tools, sufficiently high yields are in fact possible to realize large-scale integration with the requisite performance [43,44].

11.5.4.3 Integration with SiGe electronic-integrated circuits

Integration of an InP PIC with a high-speed modulator driver and TIA can be achieved through assembly on a shared carrier and by connecting the modulator driver IC outputs to the PIC RF modulator sections and high-speed PDs to IC TIA inputs through wire bonds or flip-chip bonded transmission lines, respectively [45]. Electrical power and DC controls are supplied to the PIC through electrical probing of traces on the common carrier, for example, through wire bonding. Such a platform is used at Infinera to enable validation of TOSNR/ROSNR (transmitted/received optical sigma-to-noise ratio) performance of PICs early in the development phase and with a simplified level of assembly/integration. On the transmitter side, the analog data signal from an arbitrary waveform generator (AWG) or coherent digital signal processor DSP chip is fed to the TIA through RF probes. The transmitter optical signal is picked up from the PIC through a lensed fiber. On the receiver side, the transmitter optical signal is coupled to the PIC through a lensed fiber, and the TIA output is fed to a high-speed multichannel real-time oscilloscope or the receiver side of the DSP. Operation up to 800 Gbps per wavelength has been demonstrated [46,47], demonstrating an InP PIC and SiGe ASIC integration platform capable of power-efficient performance.

11.5.4.4 InP III–V photonics: key benefits and challenges

One of the main active building blocks of an optical transmission system is the coherent light source. Besides the power, performance, area, and cost advantage, some of the other benefits of integration of all active and passive optical components on a single InP chip are the ability to test optical device properties like laser power, laser threshold, waveguide losses, photodiode responsivity, phase/attenuation tuning coefficients, and others at the wafer level, through electrical probing of the chip. In addition, through RF electrical probes, basic dynamic properties like modulation bandwidth can also be obtained. These capabilities enable design verification and performance screening at the earliest possible level of wafer fabrication and a tight feedback loop for process control.

Wafer-level testing serves multiple purposes: design verification, fab process monitoring, as well as performance/defect screening. All these can be addressed with a single set-up of a commercial automated wafer prober with electrical/optical

probing, electrical/optical switch fabric, and a range of source measure units as well as optical power meters/spectrometers.

A similar approach can be used for fab process monitoring, metrology test structures for electrical characterization (e.g., n/p contacts, sheet resistance, trench isolation, implant isolation), as well as optical characterization (e.g., optical continuity between lasers/LEDs and photodiodes, interdie waveguides), can be included on wafers and/or PIC. It is straightforward to include a variety of PIC layouts and functional blocks on a single wafer. Standard or customized wafer-level probes and test hardware provide sourcing flexibility to accommodate different test scenarios.

Due to the air gap between the wafer and the chuck, wafer-level testing does have less efficient heat sinking, resulting in potential shifts in characteristics after singulation.

InP integration also allows for die-level testing with a minimal number of assembly steps. After singulation, PIC dies are facet AR coated, mounted on a carrier for better heat-sink and handling, and typically wire-bonded to a carrier for easier electrical probing (but this is not a requirement—probing can be performed directly on the PIC surface). In addition to what is tested at the wafer level, we now have access to edge-coupled waveguides.

Defect screening, design/performance verification, and setpoint calibration are typically done in the range of product operating temperature, which allows for easier optical and electrical probing.

Current state-of-the-art multisectional sampled grating distributed Bragg reflection lasers can be tuned to any C-band or L-band frequency. For fast and precise characterization of the tuning characteristics of these lasers, fast-sweeping control sources are used in combination with a high-speed wavemeter. The laser can be set to any frequency within C-band or L-band for subsequent characterization and calibration of other passive or active components on the PIC (MZ modulators/SOAs for transmitter lasers, hybrid multimode interferometers/PDs for receiver local oscillator). This enables early design verification, fabrication feedback, and defect screening before adding costly and time-consuming assembly steps. Implications of full optical integration at die level for reliability testing are discussed in Chapter 13.

11.5.4.5 Challenges

InP manufacturing has evolved into a mature technology extensively used in telecommunication systems since the 1970s. InP and Si IC use similar wafer fabrication processes and tools. However, Si IC fabrication benefits from more automated and advanced toolsets that serve larger markets. In addition, as described earlier, photonic device integration often requires multiple epitaxial growth and regrowth steps. As a result, InP PIC manufacturing often experiences higher random defect densities than Si ICs.

There are, however, two differences that play in favor of InP. First, since PICs use optical waveguides for routing the data signals, the inherent occupied critical area density of an InP PIC is significantly lower than for a Si IC device. These optical waveguides are a few μm wide, and the optical wavelength in the material

is around 0.5 μm. Typical killer defect (catastrophic defects that result in electrical or optical yields of zero devices and are created by process failures) size is much larger than the Si IC world due to the larger feature size of optical devices. Defects of the order of 1 μm or less are typically harmless for yields and reliability.

Manufacturing yield is a key metric for cost-effective manufacturing—this requires rigorous yield management and yield analysis methodology. The manufacturing flow consists of epitaxy, wafer fabrication, die fabrication, assembly, test (functional and DC), reliability (burn-in and screening), and full PIC performance test (DC, RF test, and link test). Even though the manufacturing processes shown in this figure are vastly different from Si IC foundry fabrication flows, wafer fabrication die yields that can be achieved in state-of-the-art InP-based fabs are now equivalent to that of Si CMOS circa the mid-1990s [43,48,49]. Fig. 11.27 compares line yields (normalized per 10 mask layers) for a Si CMOS fab circa 1990 and multiple generations of InP PICs. Like Si electronic ICs, the data show that line yields for four consecutive generations of InP-based PIC products are in the high 90% range.

Statistical process control is essential for any IC fab to improve product quality, reduce line variability and troubleshoot line issues. All manufacturing tools and processes are under constant process control by characterizing tool qualification wafers, witness wafers, or metrology structures (process control monitors—both in-PIC and in-cleave lane) on the product wafers. Performance of in-line metrology structures is compared to specific requirements against end-of-line performance. Product SPC, in tandem with documented and proven out-of-control action plans (OCAP), keeps all processes in control. Process capability data and parameter distributions are fed back to the device engineers so that they can make device improvements and help ensure that future designs make optimum use of the trade-offs between performance and yield.

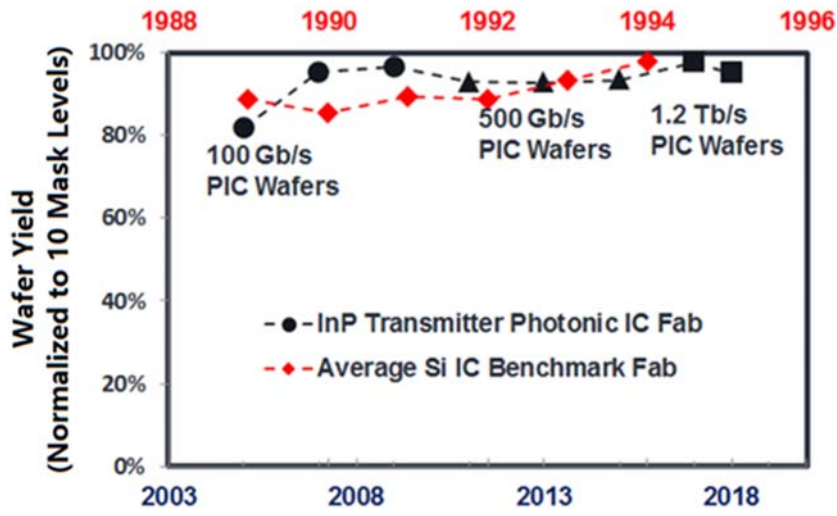

Figure 11.27 Wafer fabrication yield (normalized to ten mask levels) for multiple generations of transmitter photonic IC wafers compared to a benchmark Si CMOS fab (1989–94).

To enable high-density commercial PIC, identifying and eliminating random defects is instrumental. Extensive research conducted at Infinera helped reduce the defects, in particular the killer defects created by process failures or random defects introduced during the manufacturing process [50]. Different PIC generations required entirely different architecture, and stringent continuous improvement culture and methodology implemented in the manufacturing line reduced the defect density across multiple lines [49].

For the cost of an integrated photonic solution, it is essential to distinguish between chip cost and module cost. In some applications, the cost of the PIC is only a small part of the total cost of a module, while assembly and packaging costs are often dominant. Therefore the technology that can lead to the integration of the most functionality into a single chip will inevitably lead to the largest reduction of the assembly and packaging costs and loss in performance. Furthermore, InP fabs are often competitive to that of Si, both in terms of functionality and price. Although scaling laws for InP are similar to Si, increased wafer size and throughput can result in a lower $/mm^2 price. In general, the toolset of a Si fab is much more expensive than an InP fab and the mask costs for InP are less than silicon photonics because of the higher node typically used in the SiPh process and stricter waveguide dimensions requirements. As an example, the losses in a waveguide scale with roughness^2/width^4, in InP, this requirement is 10X less; similar scaling applies for phase errors [51]. In summary, InP-based integration offers significant performance advantages, short and medium terms, and a competitive overall optical module cost.

11.6 Photonic-integrated circuit outlook: fully monolithic photonic integration platform

One of the primary limiting factors in today's silicon photonics is the ability to increase the bandwidth, coupling efficiency, throughput, and form factor. A fully monolithic platform integrates lasers, silicon photonics components, and electronics (drivers, amplifiers, and ASIC) on the same substrate, and the same process brings the silicon photonics technology to the most advanced level. In this scheme, there is only a minimal need for packaging, and therefore parasitics associated with electronics and the limiting factors associated with extraneous coupling with photonics components can be significantly reduced. However, this scheme is complex to implement, and several fundamental challenges must be addressed. This is because different material, thermal, and process-related constraints pertain to integrating silicon electronics, III−V lasers, and silicon photonics components. Nevertheless, it is possible to create an integration scheme (e.g., III−V integration last) to achieve this goal, just as implementing copper materials to the backend of the line was achieved by incorporating several levels of contamination protocols in the manufacturing line. An integration last scheme is essential to avoid III−V contamination in the silicon transistors and optimize the time−temperature process conditions for the

III−V integration. However, environmental health and safety issues in handling the III−V compound materials in a silicon line also need to be addressed. It is much easier for a III−V platform to envisage a fully monolithic integration scheme because of its inherent material properties. Also, the speed of a transceiver circuit is dependent on the fT and fmax of the transistor. fT is limited to less than 500 GHz even in most advanced silicon technology nodes. This essentially means that technologies outside CMOS, such as HBT, HEMT structures with alternate III−V materials must be studied for improving the speed. In addition, Fully Monolithic integration has the potential to significantly lower packaging costs for both high and low volume production. Fully Monolithic integration can significantly improve these factors and thus extend the roadmap beyond 100-Gbaud NRZ data rate.

References

[1] J. Pond, G.S. C. Lamant, and R. Goldman, Software tools for integrated photonics, Optical Fiber Telecommunications. Vii, Research Gate. pp. 195−231, 2020, Available from: https://doi.org/10.1016/b978-0-12-816502-7.00005-1.

[2] S. Tanev, et al., Advances in the development of simulation tools for integrated optics devices: FDTD, BPM and mode solving techniques, in Proceedings of the Conference on Integrated Optics Devices V, San Jose, CA, January 23−25 2001, vol. 4277, Proceedings of the Society of Photo-Optical Instrumentation Engineers (Spie), 2001, pp. 1−20. Available from: https://doi.org/10.1117/12.426784. http://WOS:000169641800003.

[3] L. Chrostowski, M. Hochberg, Silicon Photonics Design: From Devices to Systems, Cambridge University Press, Cambridge, 2015.

[4] G.L. Pedrola, Beam propagation method for design of optical waveguide devices, Beam Propagation Method for Design of Optical Waveguide Devices, Wiley, 2016, pp. 1−381.

[5] D.E. McCoy, A.V. Shneidman, A.L. Davis, J. Aizenberg, Finite-difference time-domain (FDTD) optical simulations: a primer for the life sciences and bio-inspired engineering, Micron 151 (2021) 103160. Available from: https://doi.org/10.1016/j.micron.2021.103160.

[6] A.B. Fallahkhair, K.S. Li, T.E. Murphy, Vector finite difference modesolver for anisotropic dielectric waveguides, J. Light Technol. 26 (9−12) (2008) 1423−1431. Available from: https://doi.org/10.1109/jlt.2008.923643.

[7] M. Odeh, K. Twayana, K. Sloyan, J.E. Villegas, S. Chandran, M.S. Dahlem, Mode sensitivity analysis of subwavelength grating slot waveguides, IEEE Photon. J. 11 (5) (2019) 700210. Available from: https://doi.org/10.1109/jphot.2019.2939088.

[8] A. Taha, et al., Compact MMI-based AWGs in a scalable monolithic silicon photonics platform, IEEE Photon. J. 13 (4) (2021) 6. Available from: https://doi.org/10.1109/jphot.2021.3099436.

[9] I. Bork, A. Tritchkov, S.M. Shang, E. Levine, M. Zuo, MRC for curvilinear mask shapes, in: Photomask Technology Conference, Electr Network, September 21−25, 2020, vol. 11518, Proceedings of SPIE, 2020. Available from: https://doi.org/10.1117/12.2575474. <http://WOS:000632703600015>.

[10] V.W. Guo, et al., Lithographic benefits and mask manufacturability study of curvilinear masks, in: Proceedings of SPIE Photomask Technology Conference, Monterey, CA, September 17−19, 2018, vol. 10810, 2018. Available from: https://doi.org/10.1117/12.2501973. <http://WOS:000454575200015>.

[11] C. Meagher et al., Patterning challenges for monolithic silicon photonics, in: 29th Annual SEMI Advanced Semiconductor Manufacturing Conference (ASMC), Saratoga Springs, NY, April 30−May 03, 2018, Advanced Semiconductor Manufacturing Conference and Workshop-Proceedings, 2018, pp. 155−158. <http://WOS:000448893500035>.
[12] H. Choi, W.L. Wang, C. Kallingal, Quantitative evaluation of manufacturability and performance for ILT produced mask shapes using a single objective function, in: Proceedings of SPIE Conference on Design-Process-Technology Co-Optimization for Manufacturability IX, San Jose, CA, February 25−26 2015, vol. 9427, 2015. Available from: https://doi.org/10.1117/12.2087111. <http://WOS:000354257500016>.
[13] Y.S. Bian, et al., Monolithically integrated silicon nitride platform, in: Proceedings of Optical Fiber Communications Conference and Exhibition (OFC), Electr Network, June 06−11, 2021, New York: IEEE, 2021.
[14] S. Chandran, et al., Beam shaping for ultra-compact waveguide crossings on monolithic silicon photonics platform, Opt.Lett. 45 (22) (2020) 6230−6233. Available from: https://doi.org/10.1364/ol.402446.
[15] F. Van Laere, et al., Compact and highly efficient grating couplers between optical fiber and nanophotonic waveguides, J. Light. Technol. 25 (1) (2007) 151−156. Available from: https://doi.org/10.1109/jlt.2006.888164.
[16] L. Bogaerts, et al., High-speed TSV integration in an active silicon photonics interposer platform, in: Proceedings of the SOI-3D-Subthreshold Microelectronics Technology Unified Conference (S3S), Burlingame, CA, 2018, IEEE SOI-3D-Subthreshold Microelectronics Technology Unified Conference, 2018.
[17] C. Okoro, et al., Elimination of the axial deformation problem of Cu-TSV in 3D integration, in: Proceedings of the Eleventh International Workshop on Stress-Induced Phenomena in Metallization, Bad Schandau, Germany, April 12−14 2010, vol. 1300, AIP Conference Proceedings, 2010, p. 214.
[18] A. Peczek, C. Mai, G. Winzer, L. Zimmermann, Comparison of cutback method and optical backscatter reflectometry for wafer level waveguide characterization, in: Proceedings of the IEEE Thirty-third International Conference on Microelectronic Test Structures (ICMTS), Univ. of Edinburgh, Edinburgh, UK, May 04−18, 2020, IEEE International Conference on Microelectronic Test Structures (ICMTS), pp. 165−167.
[19] A.P. Jacob, T. Letavic, Silicon photonics: 300mm foundry industrialization, in: Proceedings of the IEEE Optical Interconnect Conference (OIC), Santa Fe, New Mexico, 2017.
[20] Y. Yang, et al., 3D silicon photonics packaging based on TSV interposer for high density on-board optics module, in: Proceedings of the Sixty-Sixth IEEE Electronic Components and Technology Conference (ECTC), Las Vegas, NV, May 31−June 03, 2016, Electronic Components and Technology Conference, 2016, pp. 483−489. Available from: https://doi.org/10.1109/ectc.2016.89. <http://WOS:000386103500074>.
[21] D. Guermandi, et al., TSV-assisted hybrid FinFET CMOS—silicon photonics technology for high density optical I/O, in: Proceedings of the Forty-Fourth European Conference on Optical Communication (ECOC 2019), September 22−26, 2019, pp. 1−4. Available from: https://doi.org/10.1049/cp.2019.0990.
[22] M. Rakowski, et al., Hybrid 14nm FinFET—silicon photonics technology for low-power Tb/s/mm(2) optical I/O, in: Proceedings of the IEEE Symposium on Vlsi Technology, pp. 221−222, 2018.
[23] K. Miyaguchi, et al., 110 GHz through-silicon via's integrated in silicon photonics interposers for next-generation optical modules, in: Proceedings of the European

Conference on Optical Communication (ECOC), September 13–16, 2021, pp. 1–4. Available from: https://doi.org/10.1109/ECOC52684.2021.9605927.
[24] A.W. Fang, H. Park, O. Cohen, R. Jones, M.J. Paniccia, J.E. Bowers, Electrically pumped hybrid AlGaInAs-silicon evanescent laser, Opt. Express 14 (20) (2006) 9203–9210. Available from: https://doi.org/10.1364/oe.14.009203.
[25] R. Blum, Integrated silicon photonics for high-volume data center applications, in: Proceedings of the Conference on Optical Interconnects XX, San Francisco, CA, February 04–06, 2020, vol. 11286, SPIE, 2020. Available from: https://doi.org/10.1117/12.2550326. <http://WOS:000555680600017>.
[26] A.W. Fang, et al., Hybrid silicon evanescent devices, Mater. Today 10 (7–8) (2007) 28–35. Available from: https://doi.org/10.1016/s1369-7021(07)70177-3.
[27] P. Doussiere, IEEE, Laser integration on silicon, in: Proceedings of the Fourth IEEE International Conference on Group IV Photonics (GFP), Berlin, Germany, August 23–25, 2017, IEEE International Conference on Group IV Photonics, New York: IEEE, pp. 164–165. <http://WOS:000427282500082>.
[28] R. Herrick, Reliability qualification and failure mechanisms for semiconductor lasers and fiber optic transceivers, in: Proceedings of the Optical Fiber Communication Conference: Optica Publishing Group, 2019.
[29] K. Giewont, et al., 300-mm monolithic silicon photonics foundry technology, IEEE J. Sel. Top. Quantum Electron. 25 (5) (2019) 11. Available from: https://doi.org/10.1109/jstqe.2019.2908790.
[30] D. Coenen, et al., Thermal modelling of silicon photonic ring modulator with substrate undercut, J. Light. Technol. (2022) 1. Available from: https://doi.org/10.1109/JLT.2022.3162987.
[31] M. Rakowski, et al., 45nm CMOS—silicon photonics monolithic technology (45CLO) for next-generation, low power and high speed optical interconnects, in: Proceedings of the Optical Fiber Communications Conference and Exposition (OFC), San Diego, CA, March 08–12, 2020, New York: IEEE. <http://WOS:000676346200571>.
[32] A. Borjon, et al., Process window OPC verification: dry versus immersion lithography for the 65 nm node, in: Proceedings of the Conference on Optical Microlithography XIX, San Jose, CA, February 21–24, 2006, vol. 6154, SPIE, pp. U2473–U2483. Available from: https://doi.org/10.1117/12.657056. <http://WOS:000238634000150>.
[33] Y. Bian, et al., Towards low-loss monolithic silicon and nitride photonic building blocks in state-of-the-art 300mm CMOS foundry, presented at the Frontiers in Optics 2020, 2020, Paper FW5D.2. <https://doi.org/10.1364/FIO.2020.FW5D.2>.
[34] B. Peng, et al., A CMOS compatible monolithic fiber attach solution with reliable performance and self-alignment, in: Proceedings of the Optical Fiber Communications Conference and Exposition (OFC), 2020.
[35] Y. Bian et al., 3D Integrated Laser Attach Technology on a 300-mm Monolithic CMOS Silicon Photonics Platform, IEEE J. Sel. Topics Quantum Electron., Invited Paper, vol. 29, no. 3, May/June. 2023, Art. no. 8200519. Available from: https://doi.org/10.1109/JSTQE.2023.3238290
[36] S.S. Li, High-speed III-V semiconductor devices, Semiconductor Physical Electronics, Springer, New York, NY, 2006, pp. 613–663.
[37] I. Vurgaftman, J.R. Meyer, L.R. Ram-Mohan, Band parameters for III-V compound semiconductors and their alloys, J. Appl. Phys. 89 (11) (2001) 5815–5875. Available from: https://doi.org/10.1063/1.1368156.

[38] Y.D. Galeuchet, P. Roentgen, Selective area MOVPE Of GaInAs/InP heterostructures on masked and nonplanar (100) and (111) substrates, J. Cryst. Growth 107 (1−4) (1991) 147−150. Available from: https://doi.org/10.1016/0022-0248(91)90447-d.
[39] M. Sugiyama, IEEE, Selective are growth of III-V semiconductors: From fundamental aspects to device structures, in: Proceedings of the Twenty-Second International Conference on Indium Phosphide and Related Materials, Kagawa, Japan, May 31−June 04, International Conference on Indium Phosphide and Related Materials, 2010.
[40] Y. Abe, K. Kishino, Y. Suematsu, S. Arai, GaInAsP InP Integrated laser with butt-jointed built-in Distributed-Bragg-Reflection waveguide, Electron. Lett. 17 (25-2) (1981) 945−947. Available from: https://doi.org/10.1049/el:19810661.
[41] W.D. Laidig, et al., Disorder of an AlAs-GaAs superlattice by impurity diffusion, Appl. Phys. Lett. 38 (10) (1981) 776−778. Available from: https://doi.org/10.1063/1.92159.
[42] L.A. Coldren, et al., High performance InP-based photonic ICs—a tutorial, J. Light. Technol. 29 (4) (2011) 554−570. Available from: https://doi.org/10.1109/jlt.2010.2100807.
[43] R. Nagarajan, C. Doerr, F. Kish, I.P. Kaminow, T. Li, A.E. Willner, Chapter 2—semiconductor photonic integrated circuit transmitters and receivers, Optical Fiber Telecommunications, sixth edition, Academic Press, Boston, 2013, pp. 25−98.
[44] F. Kish, et al., System-on-chip photonic integrated circuits, IEEE J. Sel. Top. Quantum Electron. 24 (1) (2018) 6100120. Available from: https://doi.org/10.1109/jstqe.2017.2717863.
[45] V. Lal, et al., 1.6Tbps coherent 2-channel transceiver using a monolithic Tx/Rx InP PIC and single SiGe ASIC, in: Proceedings of the Optical Fiber Communication Conference Exposition (OFC), 2020.
[46] R. Going, et al., 4 x 600 Gb/s photonic IC transmitter and receiver modules, in: Proceedings of the European Conference on Optical Communication (ECOC), 2018.
[47] R. Going, et al., InP-based coherent PICs for 100 Gbaud operation, in: OSA Advanced Photonics Congress (AP) 2019 (IPR, Networks, NOMA, SPPCom, PVLED), Burlingame, California, 2019: Optica Publishing Group, OSA Technical Digest, p. SpM4E.3. Available from: https://doi.org/10.1364/SPPCOM.2019.SpM4E.3. <http://opg.optica.org/abstract.cfm?URI = SPPCom-2019-SpM4E.3>.
[48] R.C. Leachman, D.A. Hodges, Benchmarking semiconductor manufacturing, IEEE Trans. Semicond. Manuf. 9 (2) (1996) 158−169. Available from: https://doi.org/10.1109/66.492810.
[49] G.E. Hoefler, et al., Foundry development of system-on-chip InP-based photonic integrated circuits, IEEE J. Sel. Top. Quantum Electron. 25 (5) (2019). Available from: https://doi.org/10.1109/jstqe.2019.2906270.
[50] F.A. Kish, et al., Current status of large-scale inp photonic integrated circuits, IEEE J. Sel. Top. Quantum Electron. 17 (6) (2011) 1470−1489. Available from: https://doi.org/10.1109/jstqe.2011.2114873.
[51] M. Smit, K. Williams, J. van der Tol, Past, present, and future of InP-based photonic integration, APL Photon. 4 (5) (2019) 050901. Available from: https://doi.org/10.1063/1.5087862.

Packaging and test technologies

12

Peter O'Brien[1], *Kamil Gradkowski*[1], *Padraic E. Morrissey*[1], *Sylwester Latkowski*[2], *Helge Gehring*[3], *Wolfram Pernice*[3,4] *and Jörg Smolenski*[5]

[1]Tyndall National Institute, Cork, Ireland, [2]Eindhoven University of Technology, Eindhoven, The Netherlands, [3]Westfälische Wilhelms-Universität Münster, Münster, Germany, [4]Heidelberg University, Heidelberg, Germany, [5]Nanoscribe GmbH & Co. KG, Eggenstein-Leopoldshafen, Germany

12.1 Introduction

The previous chapters have described many investments in programs to advance PIC technologies, resulting in the rapid growth of highly innovative PIC-based solutions for a wide range of applications and markets. However, much of this activity has focused on the PIC chip, and there are assembly and packaging bottlenecks which impede users from commercializing full system solutions. Current packaging techniques typically rely on package-by-package or "package-level" processes, where optical and electrical interconnects are completed when the photonic device has already been placed in the mechanical package. This serial style process flow is difficult to automate and has limited throughput, with scale-up in production directly related to the number of packaging machines used. Historically, such serial processes were acceptable for low-volume high-value applications, such as long-haul fiber optic communications but for emerging high-volume applications in data center and HPC markets, more scalable solutions are required. A typical serial packaging process for a telecom component involves the photonic device being assembled in the package, electrical connections formed using gold ribbon wire bonding, active alignment of the optical fiber to PIC chip to maximize coupling efficiency and hermetic sealing of the mechanical housing in an inert atmosphere. This package-level process is slow, as it typically takes tens of minutes to complete. The process is also incapable of meeting the demands of diverse applications due to the relative inflexibility of the package design and the inability to make significant changes in the process flow. PIC testing also follows the package-level approach which severely limits throughput. This situation is further complicated by the introduction of complex PIC chips which have many optical and electrical connections (I/Os), often requiring specialized electrical interconnects for high-frequency operation. The packages also have demanding thermal requirements because of the exceptionally high levels of component integration and heat generated by active

Integrated Photonics for Data Communication Applications. DOI: https://doi.org/10.1016/B978-0-323-91224-2.00007-2
© 2023 Elsevier Inc. All rights reserved.

photonic and electronic devices. Fig. 12.1 provides a high-level overview of the main optical, electrical, thermal and mechanical packaging technologies. All of these packaging technologies are usually simultaneously required for any specific application. The result is that PIC module integration, packaging and test production costs still typically account for over 50% of the overall PIC module manufacturing cost, compared to 10%–20% for microelectronic devices [1]. There is a clear need for cost-effective and scalable PIC packaging and test technology, a technology that can be easily accessed by companies using PICs and which can be integrated into their supply chains. The need for more automated packaging technologies is enormous; enhancing opportunities in existing markets and supporting the growth of new markets and applications.

12.2 PIC packaging technologies

PIC packaging is a highly interdisciplinary endeavor, requiring deep expertise in advanced packaging materials, photonic and electronic device integration processes, precision optical assembly, high-speed electrical design, thermal management, reliable mechanical design and an understanding of end-user applications to ensure package designs are fit-for-purpose. An effective PIC testing procedure must also be compatible with the packaging technology, especially package design, ensuring an efficient high-throughput test capability. Fig. 12.2 provides a summary of the key attributes and performance requirements essential for an effective PIC packaging and test technology.

12.3 Transceiver package design

As described in Chapter 1, optical transceivers are highly compact packaged modules containing active and passive photonic components, including electrical ICs,

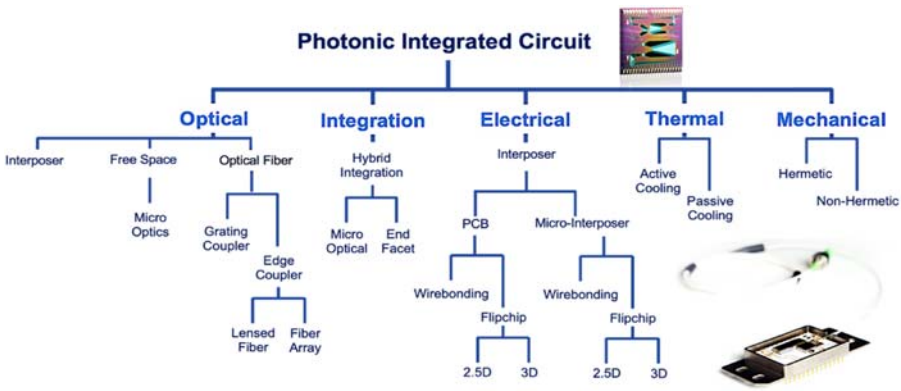

Figure 12.1 Overview of PIC packaging technologies.

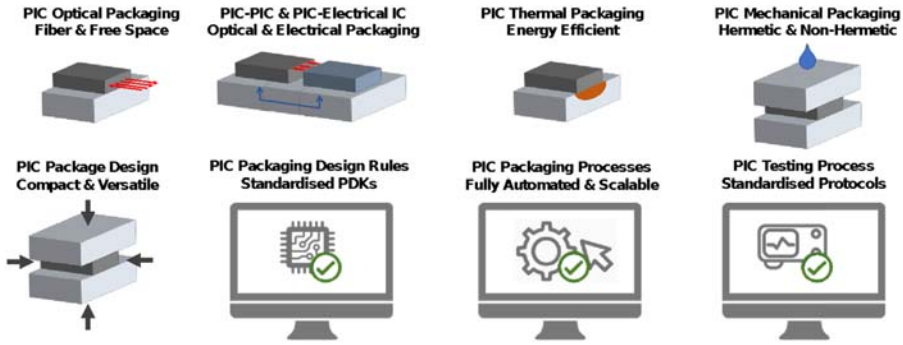

Figure 12.2 PIC packaging technology should consider all aspects of packaging and test, including; design standards; optical-electrical packaging processes; device integration; thermal management and reliability.

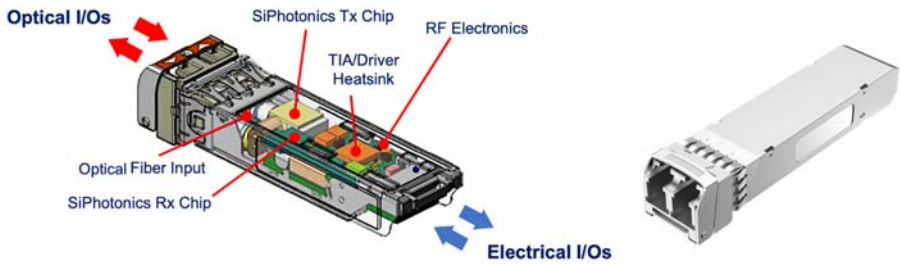

Figure 12.3 Schematics showing the configuration and main components in a standard pluggable optical transceiver module.

such as amplifiers and drivers, PCBs to support both photonic and electronic devices and to transmit electrical signals. The package also contains an array of mechanical structures to enable alignment of input/output optical fibers, dissipate heat from active components and to provide physical and electrical shielding from the outside environment. Fig. 12.3 shows a typical transceiver module design, taking the standardized SFP or Small Form Pluggable format. Variations on this design include the QSFP Small Form Pluggable format. High-speed optical signals are transmitted and received from one side of the module and electrical signals from the opposite side. The primary packaging challenge is to combine all these complex optical, electrical and mechanical functionalities in a highly compact and low-cost module. Key challenges include efficient optical coupling between single-mode optical fiber arrays (FAs) and the photonic chips, integration of a laser source with the photonic chips, integration of electronic ICs with the photonic device and routing of the high-speed electrical signals between internal devices and external system connectors.

The requirement to reduce energy consumption and increase data rates within data centers is driving progress on co-packaging of the optical module with the high-speed electrical switches, see Fig. 12.4. Locating or co-packaging the optical transceiver near the switch component reduces energy loss over copper traces on the switch PCB. This type of packaging arrangement is typically known as 2.5D integration and requires advanced electrical interposers to facilitate high-speed signal transmission and efficient heat dissipation.

12.4 Optical packaging

Current state-of-the-art photonic packaging requires permanent bonding of the FA to the photonic chip using UV- or heat-curable epoxy. This is a labor-intensive process that requires a high precision of alignment and long-term mechanical and thermal stability. Additionally, the rising requirement in the industry is reflow-compatibility of the package, which means that the bond has to survive with minimal shift through a short-term exposure to temperatures as high as 260°C. The precision that is necessary for optimal optical coupling efficiency (CE) is related to the mode-field diameter (MFD) of the coupler on the photonic device and the fiber. The larger the MFD, the more relaxed the alignment tolerances are, as shown in Fig. 12.5. From this perspective, out of the two optical main coupling schemes in existence: edge- or grating-coupling, it is the latter, which is easier to satisfy, since the grating couplers can be designed for any mode size, hence their design is typically for the single-mode fiber with MFD of around 10 μm, depending on the wavelength [2]. For edge coupling, the 3-μm MFD is typical for many Si foundries,

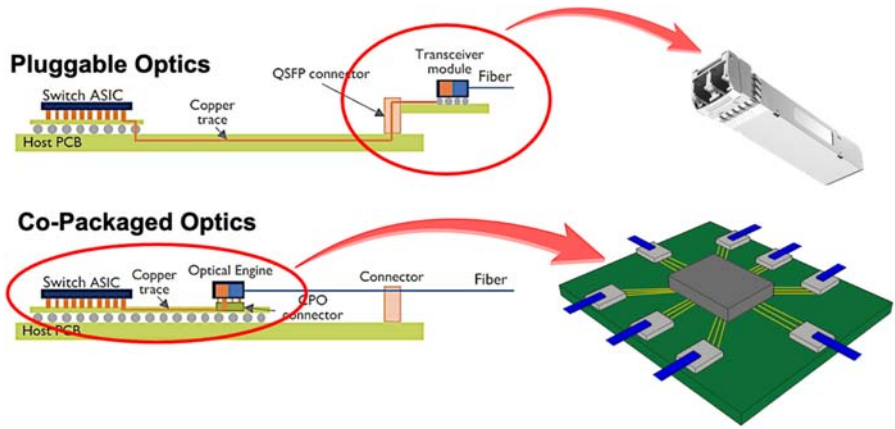

Figure 12.4 Schematics showing the evolution from pluggable optical transceivers to co-packaged systems. The drive is to move the transceiver closer to the switch component using electrical interposers.

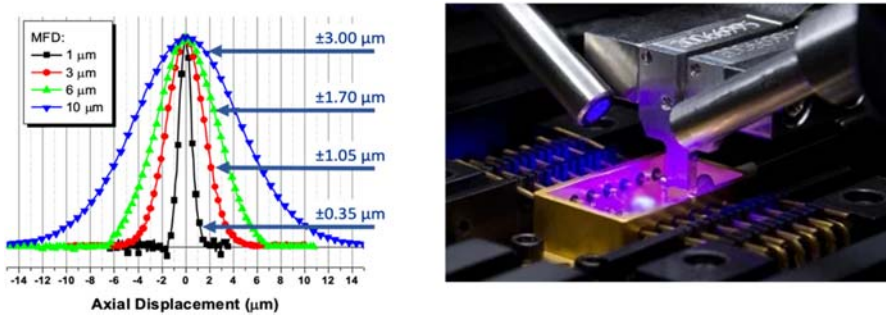

Figure 12.5 Photonic devices often have different mode sizes, ranging from 1 to 10 μm. The left image shows the alignment tolerance for 1-dB loss for various mode sizes. The right image shows UV curing of optical fibers in a photonic package.
Source: Images: Tyndall & ficonTEC.

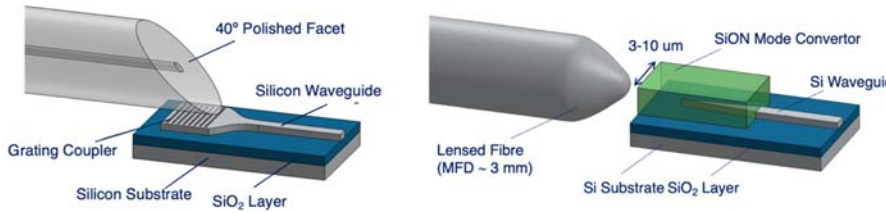

Figure 12.6 Fiber coupling to silicon photonic waveguides can be achieved using a surface grating or edge couplers. Edge coupling is becoming more common as it offers board spectral operation and mode size converters can be fabricated on the photonic devices to match the MFD of the coupling fibers.

however a number of foundries are starting to produce devices that support 10 μm mode size. An overview of both processes is shown in Fig. 12.6.

Surface grating couplers have dominated optical packaging processes, especially for silicon photonic devices [3]. They enable devices to be singulated without the need for facet polishing and can be easily imaged by packaging equipment for fast and automated packaging routines. However, grating couplers have spectral bandwidth limitations, polarization sensitivities and require the use of special fibers, such as angle-polished fibers which can be expensive. Edge coupling is becoming more common as it offers broad spectral operation and mode size converters can be fabricated on the photonic devices to match the MFD of the coupling fibers to increase coupling efficiency [4]. However, edge coupling can be a slow process, with challenges to achieve "first light," which is the process to reach first coupling between the optical fiber and photonic device. Furthermore, edge coupling is typically only possible for 1D linear arrays (both on the fiber and on the PIC) which limits the achievable "shoreline density" of optical outputs per device facet. Innovative approaches to address alignment challenges with edge coupling include

the use of integrated v-grooves for passive alignment, such as the design now offered by Global Foundries.

Typical optical packaging processes use an active alignment procedure, where the coupled optical power is continuously measured and maximized, and the fiber is fixed in position. These processes are performed at the package level and suffer from significant throughput limitations, with cycle-times on the order of minutes, depending on the complexity of the photonic device to be assembled. It should be noted that active alignment may not necessarily be the rate limiting step in this process, which may be dominated by component loading and unloading into the optical packaging machine or extended UV curing times.

An alternative to bonding the fiber directly to the photonic device involves the use of Micro Lens Arrays (MLAs). Using MLAs allows the light from a photonic device to be expanded and subsequently collimated, before it is then refocused at distance (several millimeters) into the core of the fiber using another lens [5]. The lens and the fiber at the output will be collectively called connector coupler. This principle is shown in Fig. 12.7.

The beam expansion between the photonic device and the fiber is a crucial element of the scheme because the same rule of 1 dB tolerance that applied to bare fiber with a chip, now applies to a much wider beam, see Fig. 12.8. Therefore, the

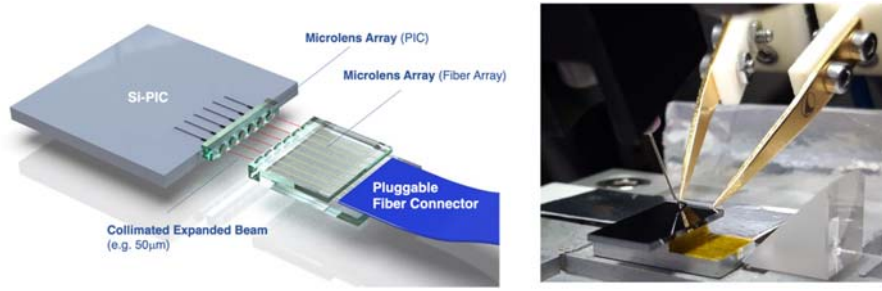

Figure 12.7 Using a microlens on the photonic device to couple to a ribbon FA (left image) and attachment of a microlens to a silicon photonic device in a packaging machine.

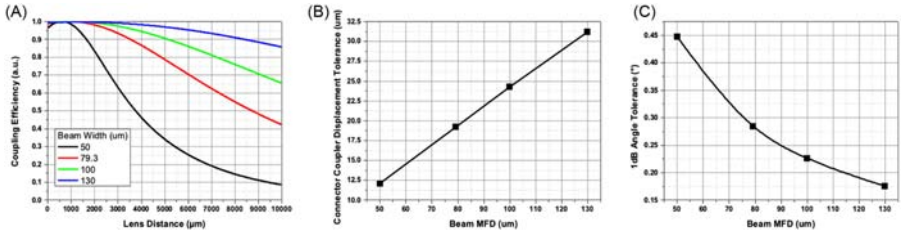

Figure 12.8 ZEMAX simulations of edge coupler expanded beam paradigm. (A) Coupling efficiency as a function of distance. Tolerance of displacement (B) and angle (C) as a function of the expanded beam diameter.

wider the beam, the more tolerant it can be for the displacement of the connector coupler. Collimation of the beam means that the beam can now only get wider through an inherent, natural divergence of a Gaussian beam. The wider the beam, the slower the divergence. This allows for placing of both sides of the link at significant distances (up to several millimeters) from one another without losing on the coupling efficiency. The optimal gap size can be tuned using design parameters of the lenses, with wider expanded beams displaying better tolerances to changes in distance and larger gap tuning range.

There are however limits to how wide the beam can be expanded. One factor that impacts this is related to manufacturing, as wider beams will require thicker lens substrates (or larger distance between the lens and the coupler), sometimes beyond what can be made on available wafers. Another factor is more pragmatic and is related to a system's performance. If the collimated beam is too wide, it can result in a cross-talk with adjacent optical channels as the light will be coupled to there as well. Therefore, the beam width (defined as a diameter where the power drops to $1/e^2$ of the maximum, center value) should not exceed around 60% of the pitch between the couplers. For a standard pitch of 250 μm this value is around 145 μm. For the 127 μm standard, the beam should be not wider than 75 μm.

Evanescent optical coupling [6] is another technique that has been shown to offer an efficient mechanism to couple light between the optical fiber and PIC waveguide via an optical interposer, much like an electrical interposer is used to interface between adjacent electronic devices. The photonic device can be structured at the wafer-level to include an evanescent coupling element, and when bonded on top of the optical interposer via a flip-chip process, enables light to efficiently transfer between the photonic device and interposer waveguides [7]. An added advantage of the evanescent technique is the ability to form an optical connection at any point on the PIC surface, rather than at the PIC edge or facet. Optical interposers can be fabricated in polymer or glass, with embedded waveguides interfacing between the optical fibers and photonic device [8]. Although at a relatively early stage of development, both micro-optics and evanescent coupling indicate the benefits of using wafer-level processes to ease the burden of optical packaging.

12.5 Electrical packaging

Electrical packaging is often overlooked but can be the largest contributor to production costs within the PIC package. This is especially the case of high-speed photonic transceivers where the package must incorporate high-speed electrical connectors, such as V, K or GPPO RF connectors. As the demand for higher operating frequencies increases, there is also a need to integrate the electronic chips closer to the PIC, known as co-packaging. Co-packaging typically requires the use of an electrical interposer. Research on interposers for PICs has picked-up in the past 5

years, with the integration of PICs and electronics on the interposer platform containing through-vias not yet widely available.

Wire bonding is the most common method of providing an electrical connection from a PIC to a PCB within a package. Au or Al ball or wedge bonding is commonly used for DC applications while ribbon wire bonds are used for high-speed signals (RF). An example of typical electrical wire bonds is shown in Fig. 12.9.

Where there is a mechanical mismatch between the photonic device and PCB bond pads, electrical interposers are used as intermediate structures to bridge them together, see Fig. 12.10. These can take the form of Ceramics (AlN), Silicon or Glass interposers.

Where high frequencies are required, typically over 20 GHz, it is advisable not to use ribbon wire bonds and to revert to flipchip packaging. This requires more complex electrical interposers which include through-substrate vias (TSVs) and multilevel electrical redistribution layers (RDLs). These interposers are usually connected to the PCB motherboard using solder bump bonds on the underside of the interposer, where the solder

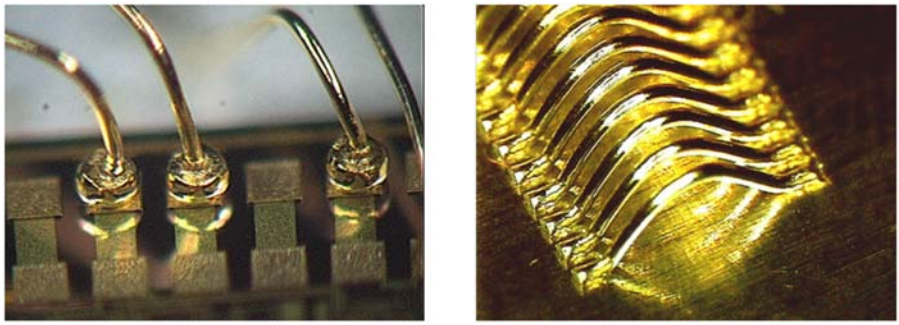

Figure 12.9 Examples of electrical wire bonds used in photonic packages. Gold ball bond (left image) and Ribbon bond (right image).

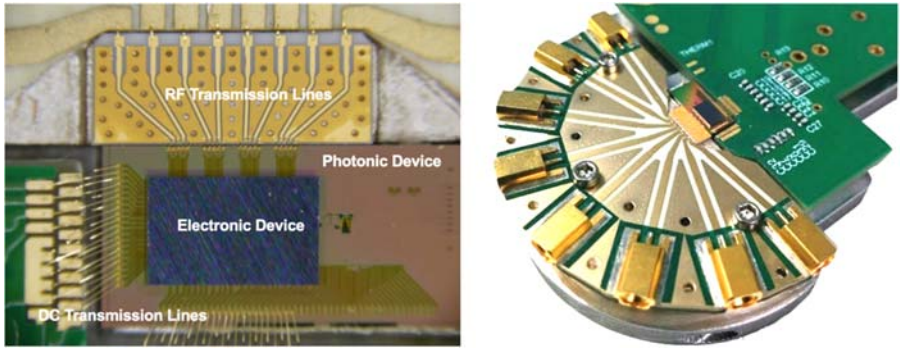

Figure 12.10 AlN ceramic interposer is used to fan-out RF transmission lines from the photonic device to the larger pitched PCB (left image). The larger pitch PCB transmits the RF lines to the high-speed connectors (right image).

bumps eliminate the need for wire bonding between the interposer and PCB. Electrical transitions between the photonic and electronic devices and interposers, interposer vias and interposer to PCB need to be carefully designed to minimize signal parasitic effects. The interposer material should also be carefully selected to ensure maximum heat transfer from the photonic and electronic devices and to minimize mechanical stress between the devices and PCB motherboard.

An example of such a high-speed interposer for 2.5D integration is shown in Fig. 12.11. A PIC and 5 × TIAs are flip-chip bonded to a Silicon interposer. An X-ray image of the same system is also shown, and the solder connections between interposer and PIC/IC can clearly be seen. The interposer utilizes TSVs to bring the RF/DC signals from the top surface through to the bottom. Once at the bottom of the interposer, a 2D array of ball grid array (BGA) pads are used to interface with a high-speed PCB. This electrical connection is shown in Fig. 12.12. The PCB fan-outs these BGA connections

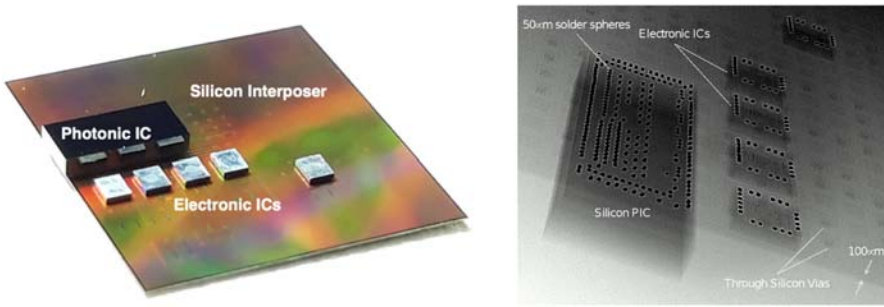

Figure 12.11 Co-packaging of photonic device and electronic ICs on a silicon electrical interposer (left image). X-ray image of the same package showing the solder bumps and through silicon vias connection to the underside of the interposer. This enables the assembly to be directly placed and soldered onto a PCB motherboard (right image).

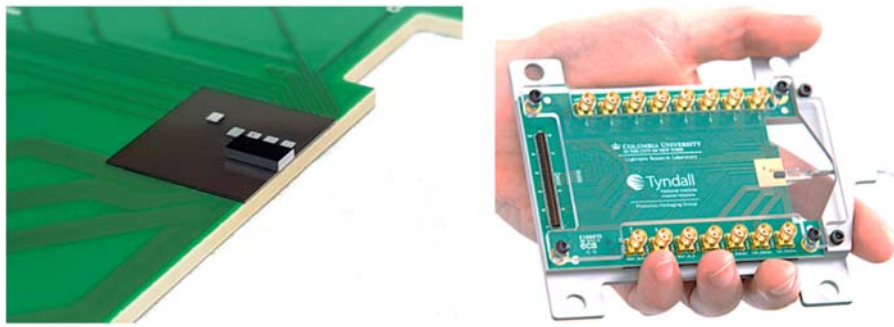

Figure 12.12 (Left) 2.5D packaging of a PIC with multiple TIAs on a silicon interposer with TSVs. (Right) The complete photonic package. A high-speed PCB fans out electrical connections to RF and DC connectors at the edges of the board. One side of the package is reserved for optical interfacing.

to a series of high-speed and DC connectors on the edge of the board. One side of the package is reserved for the optical connection and the electrical interfaces need to be codesigned with this in mind.

12.6 Hybrid laser-to-PIC packaging

As described in Chapter 2, PICs, in particular silicon and silicon nitride (SiN) devices, do not have an intrinsic light source, such as a laser or optical amplifier. There are several approaches under development to integrate light sources with PIC devices. Light source integration generally falls into three distinct approaches:

1. Off-package light source with fiber coupling between the light source and PIC waveguide. This approach is commonly used in many datacom applications as it ensures the light source and PIC are independent in terms of reliability and thermal management, where the integration process is based on conventional fiber coupling technology. However, these processes are relatively slow and expensive, resulting in a large module footprint.
2. In-package light source assembly or Hybrid Integration, which usually involves micro-optics to couple the light source and PIC [9]. Although the process is more technically challenging than fiber coupling, it enables a more compact module, while still minimizing thermal issues as the devices can be separated to avoid thermal crosstalk. Another advantage of this approach is the ability to select known good devices, ensuring process yields can be maximized. However, major challenges remain to develop a fast assembly process that maximize coupling efficiency within a small package footprint.
3. Direct bonding of the light source on the PIC chip itself, often known as Heterogeneous Integration [10]. Processes such as transfer printing enable the fast integration of light sources on the PIC chip, with optical coupling achieved through end-facet or evanescent coupling. A limitation of this process is the lack of device testing and preselection prior to the bonding process and the potential for heating the PIC device due to direct light source integration.

12.7 Packaging design rules and standardization

The need to create standardized design rules for photonic packaging is slowly gaining recognition in the PIC community, driven by the dual advantages of removing a design-burden from researchers/engineers, and streamlining the design-to-device process. Developers of photonic packaging solutions are working with industry partners to establish packaging/assembly design kits (ADKs) and rules (PDRs) to help users ensure that their Si-PICs and photonic devices are compatible with reliable and cost-effective best-practices in photonic packaging The implementation of these packaging design rules will be especially beneficial to new entrants of the Si-photonics space, by helping avoid costly design mistakes that could result in a newly designed Si-PIC being impossible to fully package. Critical PDR parameters include the pitch of grating-arrays on the Si-PIC, the pitch of fibers in the fiber-arrays, the pitch of DC and RF bond-pads, the metallization-type of bond-pads, and

the location of coupling elements (grating couplers and edge-couples) with respect to the electrical interconnects. Although currently at an early stage of development, it is expected that these packaging design rules will become more widely adopted, especially at the layout phase of the chip-design process. To increase their ease of use, the ADKs are being implemented in software layout tools and templates. This will allow designers to easily access and implement the packaging rules, in conjunction with the photonic elements available from different Si-photonic foundries.

Integration of packaging software design tools is also required in the future. As described in Chapter 10, current design process flows typically involve separate optical, electrical, thermal and mechanical phases. Future developments will require the seamless integration of these different tools, enabling designers to view how changes in one aspect of package affect others. Finally, as packaging typically consumes the majority of PIC module costs, accurate cost modeling software needs to be developed. This software tool should also be linked with design tools, providing designers with the ability to assess the cost implications of various package design changes.

12.8 Testing

Increasing maturity and complexity of Photonic Integrated Circuits (PIC) make test and characterization inevitable and frequently challenging processes associated with integrated photonics-based products [9,10]. Integrated photonics is advantageous over discrete components and bulk optics-based solutions by offering a large degree of integration leading to reduced footprint, improved stability, and reduced energy consumption. Increasing availability and accessibility to photonic integration technologies via open access foundry services that are based on standardized, generic processes [11,12] broadens a spectrum of applications for photonic ICs. A number of photonic integration technology platforms using different material systems is available in photonic foundries [13−18]. Such photonic foundries call for the emergence of other services at different points of the PIC manufacturing chain, as schematically depicted in Fig. 12.13A.

A variety of electronic-photonic design automation (EPDA) tools are available from several vendors [20−26]. Software packages support creation of layout and provide tools for simulation at the component and circuit level. The availability and functionality of process design kits (PDKs) for the main technology platforms continuously increases. Offerings from design houses and packaging service providers complete the manufacturing chain from an idea to module, see Fig. 12.13A. The increasing maturity of the EPDA tools, PDKs and generic technology processes allows for the development of proof-of-concept demonstrators, advanced prototyping, and small-to-medium volume production. Despite those advantages, many of the technology processes and tools in the manufacturing chain of PICs require substantial improvements in terms of standardization at the design [27], the front-end

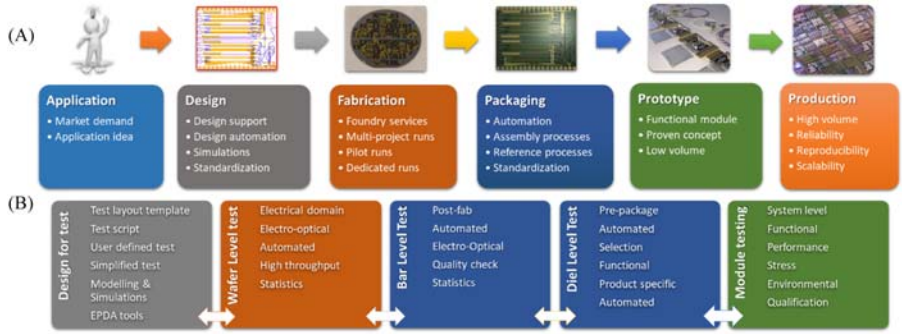

Figure 12.13 (A) Manufacturing chain of photonic integrated circuit with (B) distribution of test processes from a design to product [19].
Source: After JePPIX roadmap 2018, JePPIX. <http://www.jeppix.eu/vision/>; S. Latkowski, W. Yao, D. Pustakhod, X.J.M. Leijtens, K.A. Williams, Test methods and processes in manufacturing chain of photonic integrated circuits, in: 2018 20th International Conference on Transparent Optical Networks (ICTON), July 2018, pp. 1–4. https://doi.org/10.1109/ICTON.2018.8473652; Nazca, Nazca Design. <https://nazca-design.org/>.

fabrication [12], and the back-end processes [28]. In addition, an increased level in automation of test, assembly and packaging processes is necessary [29].

In semiconductor electronics, testing for electronic integrated circuits (EIC) is a mature process supported by practices and equipment that have been optimized to drive down the cost and time spent on testing. Development of similar methods and tools for the photonic integrated circuit is still at an early stage and the extra complexity that arises from having to measure both in the optical and the electrical domains poses extra challenges.

The term PIC refers to an immensely diverse field of different implementations where we need to consider different materials (InP, GaAs, Si, polymer, Triplex, glass), different integration schemes, such as monolithic, heterogeneous, and hybrid, assembly and packaging processes (hermetic, nonhermetic, different materials), and optical interfaces (gratings, edge, mirrors coupled to single fibers, FAs, lensed fibers, micro-optics, etc.). This diversified blend of technologies leads to a first key development area, standardization of test metrics. New standardized methodologies and qualification parameters need to be devised that apply to all technologies, types of packages, and all relevant environmental conditions, leading to a truly platform-agnostic test solution. A second area of importance is to consolidate the design and test workflow, see example in Fig. 12.13B. Variations in dimensional and physical properties of materials and modules need to be understood and considered during design. This permits engineers to predict the influence of process variations on measurement results and allows them to design dedicated and improved test structures up front. By repeating these steps in combination with a careful analysis of the stored data, the number of devices to be measured and tracked can be reduced and the functional yield can be expected to increase. This targeted reduction in the

number of devices brings us to a third key area which is test time reduction. There is a clear need for fully automated test systems. On one hand this includes inline and where possible in-situ process testing at wafer level, such as critical-dimension (CD) monitoring, defects counting, ellipsometry, etc. On the other hand, this includes the (out-of-line) automatic functional testing at wafer, bar/die and module level. For the functional test a massively parallel test approach is envisaged to bring down measurement time and cost.

More specifically for wafer-level testing, this highly parallel test approach can be enabled by scalable and modular test equipment and an increase in the number of electrical and optical input-output (IO) ports per test site. For electrical measurement instrumentation this modular approach is already quite well established; for optical instruments this is an emerging concept. To increase the number of optical IOs per test site from 10 to 100 s of couplers in the next 10 years, multicore fibers or FAs will have to be used in combination with an optical interposer to reduce the pitch of optical IOs. Measuring optical signals indirectly using on-chip photodiodes is another interesting option to eliminate the need for optical alignment.

In contrast to silicon photonic (SiPh) chips, InP- and GaAs-chips often require bar/die level testing. The main reason is that often cleaved facets with ultra-low antireflection coating are needed for full device operation. In the future the target will be to replace cleaved facets by on-wafer etched facets and to replace facet coating by on-wafer coating. Then most of the bar/die testing can be avoided and wafer level testing can be used. In addition to functional testing, the area of reliability and lifetime testing also needs to be addressed. To date, commercially available life-time testing equipment is mainly based on fiber attached mounted single dies or mounted single devices that are being tested with free space measurement setups (e.g., laser lifetime measurements using large area photodiodes). Such investigations therefore can be prohibitively expensive. Performing these tests at wafer level, simultaneously on multiple dies is an option to be investigated. This will also require developments in terms of instrumentation, for example, high-power laser sources for accelerated lifetime testing.

12.8.1 PIC testing equipment

Testing throughput is defined by the processes involved. Currently optical alignment processes between lensed fiber and a photonic IC in manual and quasi automated test setups are on the order of several minutes and tens of seconds respectively. Modern mechatronic systems with hardware assisted trajectory scanning allow for execution of such alignment in less than 2 s, therefore one or two orders of magnitude reduction in time. Automated die-tester systems enable electronic-photonic testing with part exchange, optical alignment, and electrical probing within a few seconds.

With a load of 4 trays, and a test sequence of 2 s, 100 dies can be tested within 13.5 min for a minimum set of tests. This allows more sophisticated test and measurement campaigns to be programmed with higher accuracy, reproducibility, and reliability, which is otherwise practically unachievable with manual measurement

approaches. Typical results for reproducibility of such an automated optical alignment system are show in Fig. 12.14. Generic methodologies are expected to become available in the next two years, which will remove the need for reinvention, while allowing the designer to focus on assessment strategies. Accelerated throughputs and customizations enabled by modularity of the equipment provided by vendors increases competitiveness and relevance of their products to the market. Examples of equipment for die testing with automated alignment are shown in Fig. 12.15.

Telecommunications test equipment, components, and methods were and are being adopted for optical testing of products used for nonlong-haul applications. The traditional methods are being extended and new methods developed to address test needs for photonic wafers, photonic integrated circuits, and System-in-Package (SiP) that

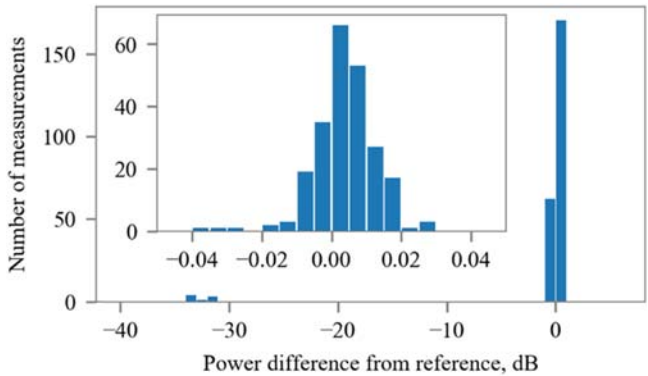

Figure 12.14 Reproducibility of automated optical alignment systems [30].
Source: After S. Latkowski, D. Pustakhod, M. Chatzimichailidis, W. Yao, X.J.M. Leijtens, Open standards for automation of testing of photonic integrated circuits, IEEE J. Sel. Top. Quantum Electron., 25 (5) (2019) 1−8. https://doi.org/10.1109/JSTQE.2019.2921401.

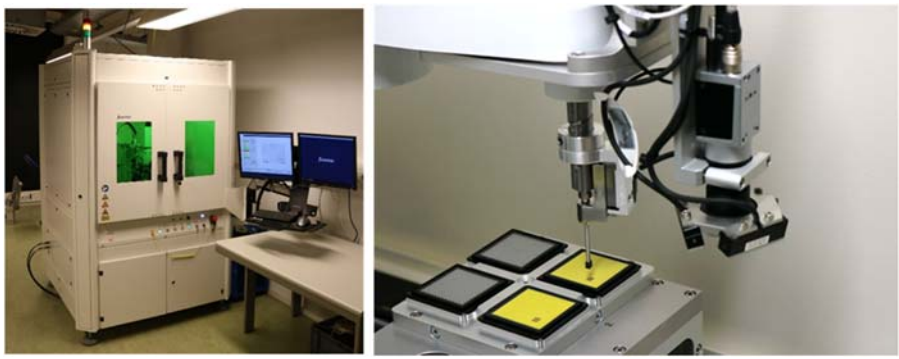

Figure 12.15 Automated die tester developed by Ficontec GmbH (left), with a modular pick and place tool for automated device under test loading and sorting (right).

utilize optics and complete systems. Utilizing these extended methods requires optical probing of both wafers and dies combined with electrical probing resulting in a series of mechanical issues. The inclusion of optical probing, especially single mode probing, requires gratings or other access points on wafers. New approaches, such as periscopes printed on single mode or polarization-maintaining FAs (periscopic lensed FAs), enable parallel and faster processing [31]. For individual die, dual mode (electrical and optical) probing is especially difficult due to the small size of die and difficulty of holding and locating probes accurately. At the SiP level, the problems are easier because the device is larger, not as fragile, and is often designed to facilitate dual media probing. The wafer, die, and SiP probe fixtures tend to be expensive due to the complexity and accuracy required. System level test access is usually easier because at that level, electrical interfaces and optical connectors are included as part of the device under test (DUT).

12.8.2 Standardized PIC testing protocols

User defined test protocols, scripting methods for controlling automated die and wafer testers and standardized data exchange are being aligned to the methods used for circuit simulation and chip design [32].

This builds design-for-test directly into the PIC product creation workflow. Automated, high throughput, modular test equipment for photonic ICs is currently being developed and a typical flow is shown in Fig. 12.16. Such modular test

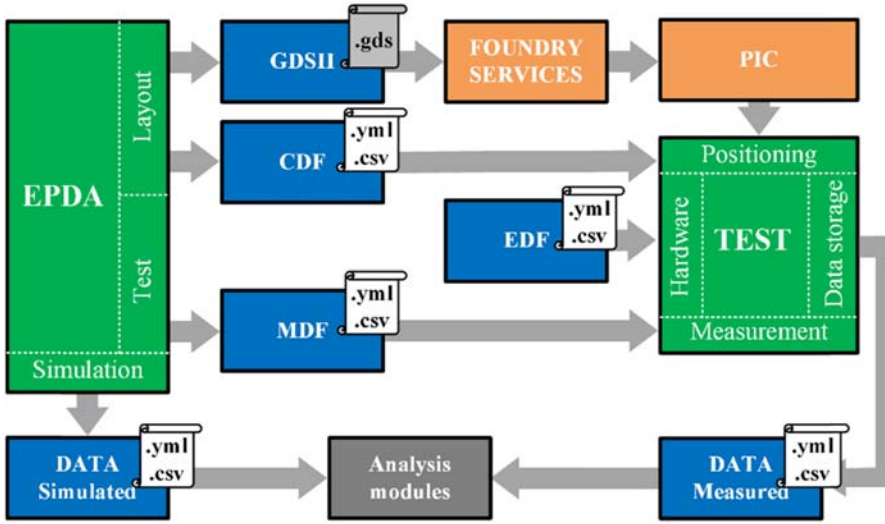

Figure 12.16 An example of the modular test framework with open standard data formats. Source: After S. Latkowski, D. Pustakhod, M. Chatzimichailidis, W. Yao, X.J.M. Leijtens, Open standards for automation of testing of photonic integrated circuits, IEEE J. Sel. Top. Quantum Electron. 25 (5) (2019) 1–8. https://doi.org/10.1109/JSTQE.2019.2921401 [30]; openEPDA Overview—openEPDATM documentation. <https://openepda.org/>.

systems are expected to support a broad range of optical, electrical and microwave test functions as well as accommodate photonic ICs fabricated in any technology platform.

The scripting methods for custom test protocols are being aligned with the design tools with an ambition to facilitate pretesting. Running such test scripts within commercial and open software tools which are supporting the PDKs provides a powerful way to introduce testing at an early phase, validate the test methods and provide direct like-for-like comparisons between the design intent and outcomes.

12.8.3 The roadmap for PIC testing

Measurement capabilities are inherent in the test automation strategies, enabling the systematic characterization of the basic building blocks, leading to detailed insights into variations. Building block characterization will be continuous, combining design intent and process control module data as well as dedicated test cells, ensuring data used for PDK tolerances will come from testing throughout the production flow, enabling statistical significance while maintaining high throughput. The first methods for automated test and recalibration were put in place in 2020, with a scale up to include critical to quality metrics by 2022 and a comprehensive recalibration regime with sustained improvements in tolerances following on [11,29,33].

Automation of test and measurement are under development as a generic designer-scripted test-automation environment [19,29,30,33]. While the precise test protocols will differ for wafer, die and module testing, the vision is to develop a library of automated test functions which can be configured to suit custom functional requirements of application specific IC at all stages of product creation according to designers' specification. This capability is envisaged as a service at a remote site or test service provider in the first instance [34]. The first use cases are anticipated for known-good-die selection, accelerating the selection of die earlier in the pilot manufacturing phases, but uses for failure mode analysis and functional testing will also be introduced over time. Businesses and innovators investing in test hardware would also be able to leverage these developments as the methods and libraries can be seamlessly deployed on their own infrastructure. Generic testing methods include building block specific attributes ranging from typical electro-optical characteristics and their dependency on the process variations to electro-static-discharge resilience, damp-heat test resilience, and lifetime attributes.

In the future, optical test methods to simultaneously characterize and compare multiple optical lanes, channels and/or ports at the same time will be needed. In particular, one need is comparative simultaneous testing of multiple signals from arrays of ribbon fibers, waveguides, sources or detectors for optical skew, jitter, etc. A related need is the ability to simultaneously evaluate optical signals multiplexed on one fiber or waveguide. Applications with arrays exceeding 256 ports (fibers or waveguides) or >256 multiplexed wavelengths are forecast in the next ten years.

In addition to the standard telecom optical parameters, such as power, wavelength, attenuation, jitter, and signal-to-noise (SNR) ratio, emerging applications aim to utilize virtually every parameter that light can have, potentially requiring the

extension of test capability in multiple dimensions, such as polarization, phase noise, amplitude noise, spatial modes, and multiple fiber cores. While these emerging needs are potentially very broad, the near-term emerging needs will most likely be driven by data communications needs.

Optical communication applications are likely to utilize 650 nm to 2000 nm wavelengths, multiplexed wavelength spacing down to 25 GHz, detector responsivity of ~1 A/W, receiver sensitivity as great as −45 dBm, power levels of 1 W or less, symbol rates of 100 GBd per lane, modulation schemes utilizing up to 10 bits per symbol, polarization multiplexing, BERs of 10−12, etc. Over time, these parameters will improve, so test capabilities will need to stay ahead of them. Data rates as high as 500 Tbps per fiber are likely to emerge in the next 10−15 years.

Regardless of application, the principles of *design for test* remain the same: The use of optical test access points, Built-In Self-Test (BIST), redundancy for self-repair, re-purpose and prognostics to report changes and deterioration during operation over the life cycle of optical products are desirable and of value in an increasing number of applications. These tests should be considered for inclusion not only in designs, but also in software design tools as well.

12.9 3D printing of micro-optics for integrated photonics

Micro optics are becoming an essential component for coupling to integrated photonic devices. Conventional micro-optic components are typically fabricated using standard semiconductor wafer-level processing techniques. However, currently integration of those micro-lenses in a device is a serial, time-consuming process and not scaled to manufacture. An exciting alternative which is gaining interest is 3D printing of micro optics [35,36] and photonic wire bonds [37−39]. This technology is an additive manufacturing process which is capable of producing extremely complex optical components with high coupling efficiencies.

12.9.1 Process for 3D printing of micro-optics

The 3D printing process is based on the principle of two-photon polymerization (2PP). Fig. 12.17 illustrates this technique. The PIC is covered locally with a monomer (photoresin), which has a special property that the polymerization chemical reaction is initiated only when the material absorbs two photons simultaneously. This occurs when a critical threshold power density is reached. The volume undergoing polymerization is called a voxel and it is a highly-localized process. The minimum size of the voxel is limited by the size of the focal spot. This process creates a series of adjacent polymerized voxels (see Fig. 12.18A) and therefore shape accuracy is limited when creating smooth surfaces, such as a curvature of a lens. However, this method can be enhanced by modulating the power of the laser as it scans across the volume, therefore creating voxels of varying sizes that can better reflect the lens contour in a process called Two-Photon Grayscale Lithography

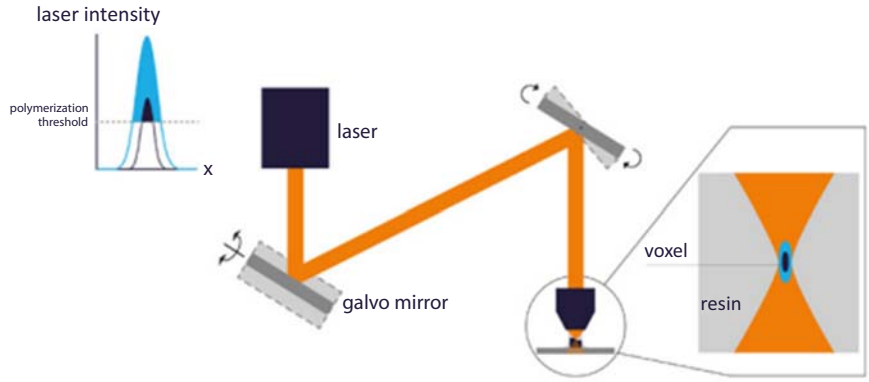

Figure 12.17 Illustration of the principle of two-photon polymerization.
Source: Nanoscribe.

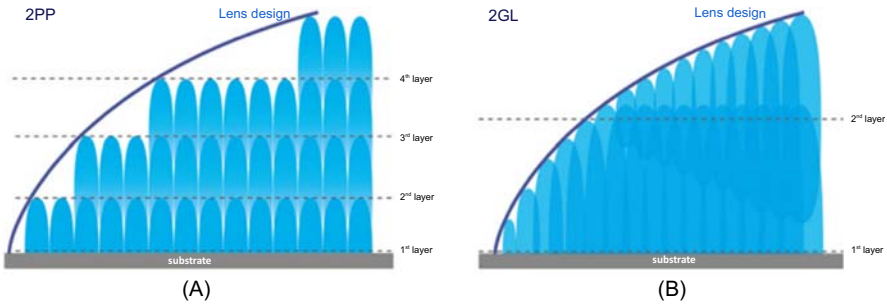

Figure 12.18 Comparison between the (A) standard 2PP and the (B) enhanced 2GL processes.
Source: Nanoscribe.

(2GL®), as is shown in Fig. 12.18B. Standard feature sizes are down to 100 nm with shape accuracy better than 200 nm and can be manufactured at a high speed. Nowadays 2PP has become a standard in the printing of optical interconnects for PIC and different 3D structures. With this 3D printing technology a wide range of materials can be used, from various photopolymers to biocompatible materials. This enables the 2PP to be utilized in a variety of applications ranging from MEMS [40], Microfluidics [41], Micro robots [42,43] up to Cell scaffolds in Bioprinting [44–46].

12.9.2 *Application of 3D printing processes for micro-optics for integrated photonics*

While the 2PP process is capable of reproducing complex shapes accurately over large wafer surfaces, it is adapted to the small and local requirements of optics on

the chips. The optomechanical set-up allows printing sizes from a few μm up to several mm. In order to achieve high expanded beam diameters, the micro-optics have to be relatively thick (e.g., 200−300 μm) in order to enable the expansion of the mode before the refraction that collimates the beam. Therefore 3D printing of free space micro optics (FSMO) is an ideal solution for direct coupling between the PIC and the fiber, as significant mode adaption can be done. It is especially advantageous for writing aspheric lenses (or any freeform optic), enabling coupling of nonsymmetric edge couplers to the circular shape of a fiber or further optical waveguides on adjacent dies. Comparisons of coupling efficiency between fiber and chip while using a lensed fiber, cleaved fiber and FSMO have shown that the lateral alignment tolerance for 1 dB excess loss can be up to 9.7 μm [36]. This enables passive alignment and pick and place in photonic packaging.

For in-plane coupling, such as in edge coupling paradigm, the lenses can be directly printed on the mode adapter at the edge of the PIC and adapt the nonsymmetric mode size for coupling into the fiber. This link can have a range of a few tens of microns up to several mm [47]. A more advanced technique can be used in applications with an angled facet, that is, when the waveguides are at a small angle with the air interface to prevent reflections back into the circuit. This method is used in InP platforms with integrated lasers. In such a case, the lenses can be printed with an integrated total internal reflection (TIR) mirror that can correct for this angle and collimate the beam at the same time for direct single-mode fiber interfacing.

For the other paradigm of out-of-plane coupling, which is currently satisfied by grating couplers, using a 3D printing process can bring the advantages of edge coupling such as a wider wavelength range and polarization insensitivity. For redirecting the light out of the plane of the chip, the microstructures can either be based on continuous bending of a polymer waveguide or on total internal reflection within the structure.

The first approach is depicted in Fig. 12.19 and starts with an adiabatic conversion of the mode from the planar waveguide to a polymer waveguide by continuously shrinking the width of the planar waveguide while enlarging the height of the polymer taper. Subsequently, the light is directed out of the plane of the chip using a continuously bent waveguide. In order to match the larger mode field size of the optical fiber, the waveguide is enlarged in the following section to increase the mode field size of the propagating beam. Finally, the structure ends with a lens to focus the light on the facet of the optical fiber. To guarantee mechanical stability, two supporting beams are attached at the front-side of the structure. As this design does not include components which are very sensitive to the wavelength of the propagating light, it allows broadband interfacing over more than all telecommunication bands [48].

For single-mode interfacing in the visible wavelength region, the structure would need to be scaled down to remain single mode in the bending region which would cause stability problems. Therefore, in the second approach, instead of relying on bent waveguides, the light is coupled from a PIC waveguide to the polymer waveguide and enlarged within the plane of the chip and subsequently redirected using a

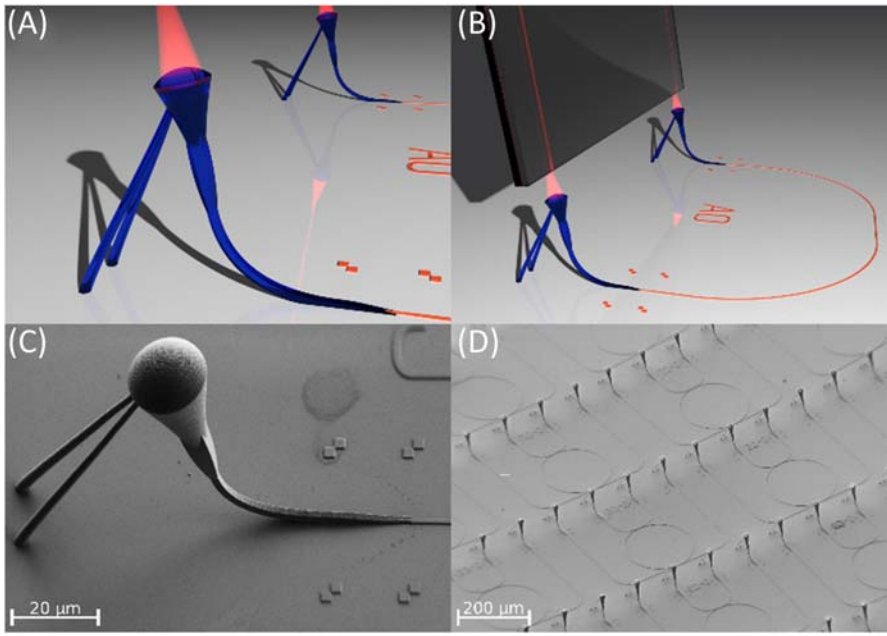

Figure 12.19 (A) Rendering of a fiber-to-chip coupling structure (blue) connected to a planar silicon nitride waveguide (orange). The focused laser beam emitted by the structure is sketched in red. (B) Two coupling structures connected by a silicon nitride waveguide and accessed by a fiber-array. (C) SEM image of the fiber-to-chip coupling structure. (D) Several coupling structures connected to integrated ring resonators [48].
Source: Reproduced from H. Gehring, M. Blaicher, W. Hartmann, P. Varytis, K. Busch, M. Wegener, et al., Low-loss fiber-to-chip couplers with ultrawide optical bandwidth, APL Photon. 4 (2019) 010801.

total reflection plane towards the optical fiber. After the total-reflection plane, the diverging beam is focused on the fiber facet using a curved surface of the coupling structure. As this structure is connected over nearly the whole length to the substrate, the coupling structure is very stable and can be scaled according to the desired wavelength. This scalable concept is presented in Fig. 12.20 [48]. A realization of such a structure for interfacing in the visible wavelength region is depicted in Fig. 12.21A in a SEM-image while in Fig. 12.21B a high number of such coupling structures is depicted in an optical microscope image, showing the possibility for many interfaces in a small region.

The transmission through this structure is depicted in Fig. 12.22, and it shows a broadband transmission over several hundreds of nanometers. This is a far larger transmission bandwidth than that of the conventional grating couplers which allow coupling over only a few tenths of nanometers. It also shows lower sensitivity to polarization.

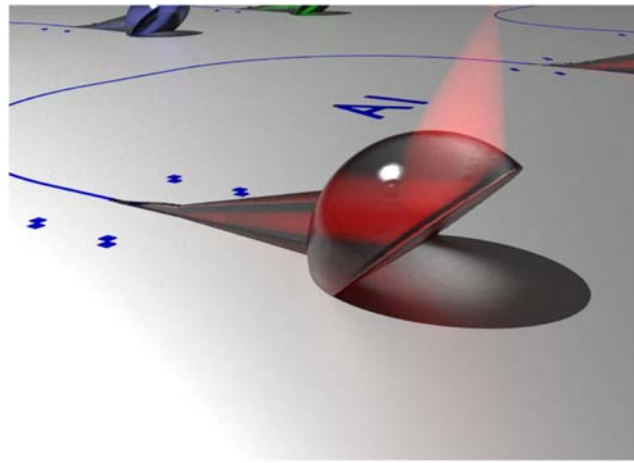

Figure 12.20 Rendering of a fiber-to-chip coupling structure (black) based on total internal reflection. The structure is connected to a planar nanophotonic waveguide (blue). The path of the light is sketched in red.

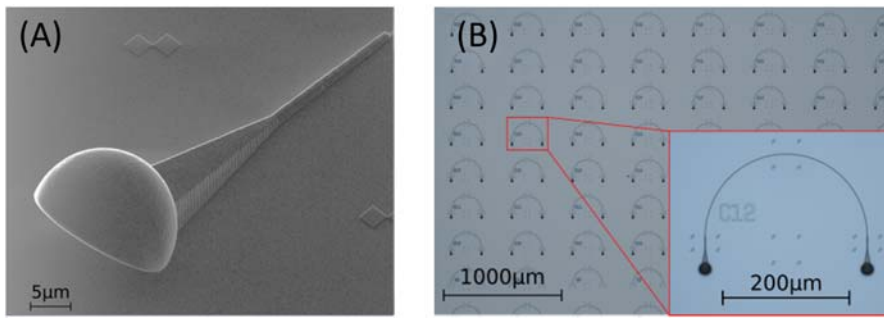

Figure 12.21 (A) SEM-micrograph of a total internal reflection coupling structure. (B) Optical micrograph of several coupling structures interconnected in pairs [49].
Source:Reproduced from H. Gehring, A. Eich, C. Schuck, W. Pernice, Broadband out-of-plane coupling at visible wavelengths, Opt. Lett. 44 (2019) 5089−5092.

As the 3D-structures are made of polymer and therefore are organic materials, the structures can be dissolved using oxygen plasma, which allows designers to realize reconfigurable nanophotonic circuits [50]. The realization of the idea is depicted in Fig. 12.23. Planar-to-polymer-waveguide interfaces are fabricated in the planar circuitry. Subsequently, the polymer waveguides are fabricated atop of the existing circuitry and allow a control of the routing between these regions. The image shows the current configuration, where a lower-left waveguide is connected to the circuit on the right, however, the trajectories of previous connections that were subsequently removed are still recognizable.

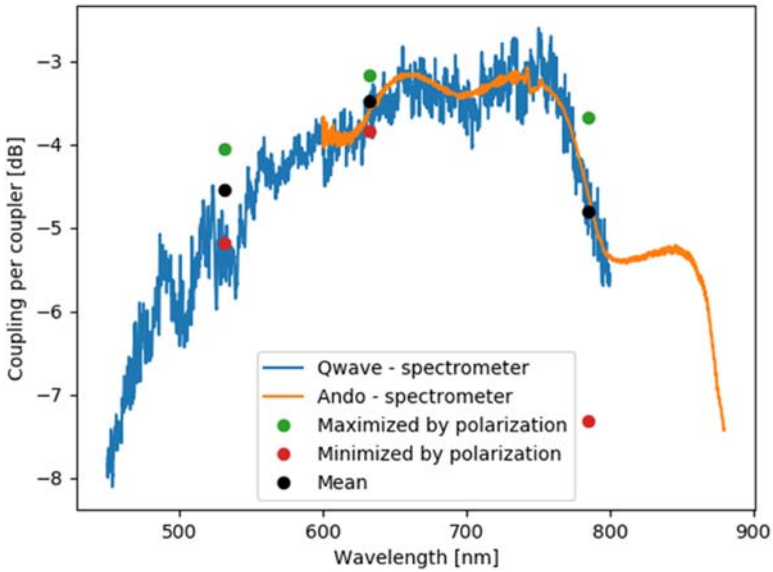

Figure 12.22 Per-coupler coupling efficiency of the total internal reflection structure. Data were obtained using an unpolarized supercontinuum source and two different spectrometers (blue/orange) and at discrete wavelength using a polarized laser source [49].
Source: Reproduced from H. Gehring, A. Eich, C. Schuck, W. Pernice, Broadband out-of-plane coupling at visible wavelengths, Opt. Lett. 44 (2019) 5089−5092.

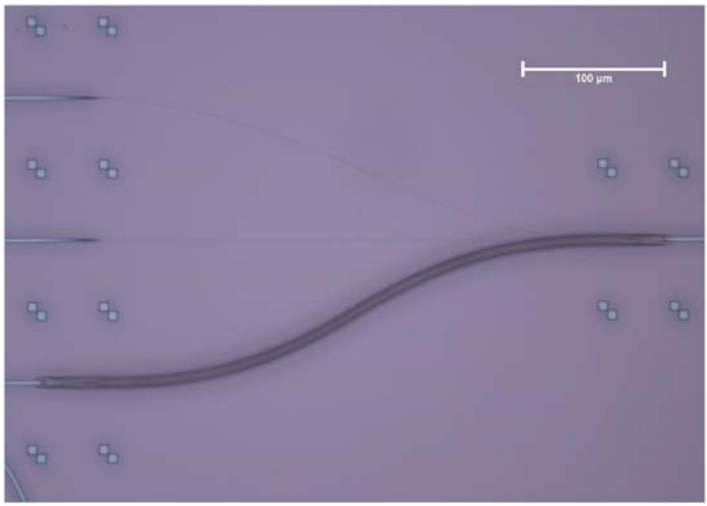

Figure 12.23 Reconfigured nanophotonic circuit. The lower-left interface is connected to the interface at the right-hand side using a polymer waveguide. The trajectory of the previous connections between the right interface and the middle and upper interface on the left-hand side are still slightly visible [51].
Source: Reproduced from C. Maibohm, O.F. Silvestre, J. Borme, M. Sinou, K. Heggarty, J.B. Nieder, Multi-beam two-photon polymerization for fast large area. 3D periodic structure fabrication for bioapplications, Sci Rep 10 (2020) 8740.

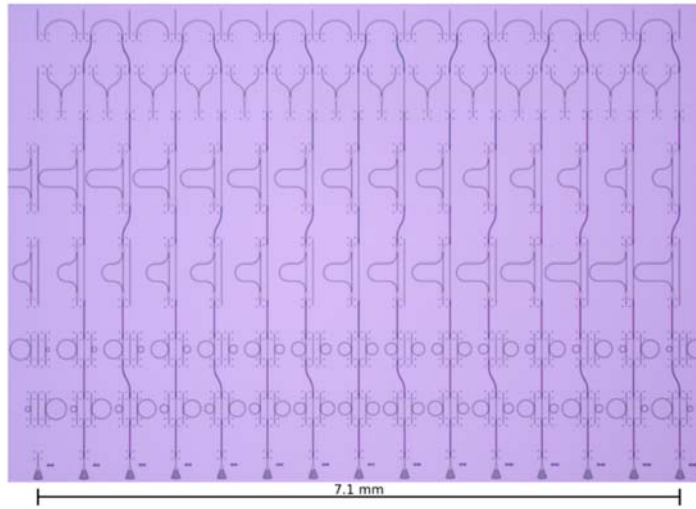

Figure 12.24 Optical micrograph of several reconfigurable nanophotonic filters consisting of elements, such as ring resonators and Mach−Zehnder interferometers. The circuit is configured using polymer waveguides fabricated using direct-laser writing [50].
Source: Reproduced from H. Gehring, M. Blaicher, T. Grottke, W.H.P. Pernice, Reconfigurable nanophotonic circuitry enabled by direct-laser-writing, IEEE J. Sel. Top. Quantum Electron. 26 (5) (2020) 1−5.

An exemplary application of this approach is presented in Fig. 12.24, where several integrated filter elements (ring resonators and Mach-Zehnder-interferometers) are fabricated on the PIC and, subsequently, are interconnected multiple times in various configurations using polymer waveguides to realize different filter-functions achieved by the planar components.

Both presented concepts for utilization of micro-optic 3D printing, that is, 3D-structures for interfacing to optical fibers and reconfigurable photonic circuits, can be combined on a single chip. This allows for an extreme degree of reconfigurability, as new on- and off-PIC optical pathways and I/Os can be written and erased.

12.9.3 The roadmap for 3D printing of photonics

The technology for 3D printing with 2PP is expected to progress rapidly for both R&D and manufacturing of integrated photonics. Based on developments, in the last years for 3D photonic wire bonds (PWB) and 3D free space micro-optics (FSMO), it is expected that in the next few years the industrial maturity of FSMO will prevail over the one of PWB. The first industrial application of FSMO for small to medium volume production is likely to appear publicly in the next few years—starting with medical sensors and instruments (e.g., OCT) and being followed by applications in the area of telecom or datacom.

On a longer time horizon, new 2PP technologies, such as multibeam based 2PP [51] or multimaterial printing [52], will make their way into various integrated photonics research labs worldwide. New materials (e.g., higher refractive index, reduced absorption, tailored transparency) will support the spreading of 3D printing over the coming 3−5 years.

12.10 Conclusion

Packaging of integrated photonic devices involves a range of complex assembly and test steps, including optical and electrical packaging, laser integration, thermal management and functional testing of the packaged module. Packaging processes typically rely on package-by-package or "package-level" processes, where optical and electrical interconnects are completed when the photonic device has already been placed in the mechanical package. This serial style process flow is difficult to automate and has limited throughput. PIC testing also follows the package-level approach which severely limits throughput. This situation is further complicated by the introduction of complex PIC chips which have many optical and electrical connections (I/Os), often requiring specialized electrical interconnects for high-frequency operation. The result is that PIC module integration, packaging and test production costs typically account for over 75% of the overall PIC module manufacturing cost, compared to 10%−20% for microelectronic devices.

The future of photonic packaging needs to move to wafer-level assembly and test processes. Package designs that will enable this transition are essential and developments in this area are starting to happen. Key in this transformation is the role of standardization of PIC layouts, electronic, thermal and mechanical designs of the package, as well as standardization of packaging processes and machines. All these considerations can be captured by the ADKs to be utilized at the very early stage of device design. New approaches to optical packaging include the use of micro-optics to enable pluggable connectors, moving away from directly bonding optical fibers to the photonic device. Further developments in this area will need to focus on surface rather than edge coupling, as the surface will provide a high density of optical I/Os and is easier to assemble due to enhanced machine vision capabilities. Electrical packaging, particularly for high frequency connections, will rely more on novel interposer materials, such as advanced ceramic, organic, and glass material, which also provide significant cost reductions for volume manufacturing. This technology has already been proven in high-volume microelectronic manufacturing, and this approach can be utilized in future photonic systems. New technologies are coming on-stream and may provide enhanced manufacturing processes. This includes additive manufacturing, such as high-resolution 3D printing of micro-optics. The challenge with these novel processes is to ensure they can be easily scaled for high-volume low-cost manufacturing.

References

[1] E.R.H. Fuchs, E.J. Bruce, R.J. Ram, F. Horst, R.E. Kirchain. Process-based cost modeling of photonics manufacture: the cost competitiveness of monolithic integration of a 1550-nm DFB laser and an electroabsorptive modulator on an InP platform, IEEE J. Light. Technol. 24 (8) 2006.

[2] R. Marchetti, C. Lacava, L. Carroll, K. Gradkowski, P. Minzioni, Coupling strategies for silicon photonics integrated chips, Photon. Res. 7 (2019) 201−239.

[3] D. Taillaert, P. Bienstman, R. Baets, Compact efficient broadband grating coupler for silicon-on-insulator waveguides, Opt. Lett. 29 (2004) 2749−2751.

[4] X. Mu, S. Wu, L. Cheng, H.Y. Fu, Edge couplers in silicon photonic integrated circuits: a review, Appl. Sci. 10 (4) (2020) 1538. Available from: https://doi.org/10.3390/app10041538.

[5] C. Scarcella, K. Gradkowski, L. Carroll, J.S. Lee, M. Duperron, D. Fowler, et al., Pluggable single-mode fiber-array-to-PIC coupling using micro-lenses, IEEE Photon. Tech. Lett. 29 (22) (2017) 1943.

[6] R. Dangel, J. Hofrichter, F. Horst, D. Jubin, A. La Porta, N. Meier, et al., Polymer waveguides for electro-optical integration in data centers and high-performance computers, Opt. Express 23 (2015) 4736−4750.

[7] A. La Porta, R. Dangel, D. Jubin, F. Horst, N. Meier, D. Chelladurai, et al., Optical coupling between polymer waveguides and a silicon photonics chip in the O-band, in: Proceedings of the Optical Fiber Communication Conference, paper M2I.2, 2016. <https://doi.org/10.1364/OFC.2016.M2I.2>.

[8] H. Hwang, P. Morrissey, J.S. Lee, J. Henriksson, T.J. Seok, M.C. Wu, et al., 128 × 128 silicon photonic mems switch package using glass interposer and pitch reducing fibre array, in: Proceedings of the IEEE Ninteenth Electronics Packaging Technology Conference, 2017, pp. 1−4.

[9] L. Carroll, et al., Photonic packaging: transforming silicon photonic integrated circuits into photonic devices, Appl. Sci. 6 (12) (2016) 426.

[10] B. Corbett, F. Atar, M. O'Brien, J. O'Callaghan, J. Justice, B. Roycroft, Heterogeneous integration of III-V based photonics with silicon, in: Novel In-Plane Semiconductor Lasers XIX, Vol. 11301: International Society for Optics and Photonics, 2020, p. 113010H.

[11] IPSR Roadmap | Knowledge and Innovation for Manufacturing Initiative. 2020. <https://photonicsmanufacturing.org>.

[12] L.M. Augustin, et al., InP-based generic foundry platform for photonic integrated circuits, IEEE J. Sel. Top. Quantum Electron. 24 (1) (2018) 6100210. Available from: https://doi.org/10.1109/JSTQE.2017.2720967.

[13] A. Rahim, T. Spuesens, R. Baets, W. Bogaerts, Open-access silicon photonics: current status and emerging initiatives, Proc. IEEE 106 (12) (2018) 2313−2330. Available from: https://doi.org/10.1109/JPROC.2018.2878686.

[14] SMART Photonics. <http://www.smartphotonics.nl/>.

[15] ePIXfab-European Silicon Photonics Alliance—European Silicon Photonics Alliance. <http://epixfab.eu/>.

[16] LioniX International MPW, *LioniX International*. <https://www.lionix-international.com/mpw-service/>, 2016.

[17] Technology Offerings, *AIM Photonics*. <https://www.aimphotonics.com/mpw>.

[18] Silicon Photonics (SiPho)—Tower Semiconductor, 2023. <https://towersemi.com/technology/rf-and-hpa/silicon-photonics-rf/>.
[19] S. Latkowski, W. Yao, D. Pustakhod, X.J.M. Leijtens, K.A. Williams, Test methods and processes in manufacturing chain of photonic integrated circuits, in: 2018 20th International Conference on Transparent Optical Networks (ICTON), July 2018, pp. 1–4. https://doi.org/10.1109/ICTON.2018.8473652.
[20] Nazca, Nazca Design, 2021, <https://nazca-design.org/>.
[21] Luceda, Luceda. 2021. <https://www.lucedaphotonics.com/en>.
[22] VPIphotonics—Application Area, Photonic Circuits. 2023. <http://vpiphotonics.com/Applications/PhotonicCircuits/>.
[23] Synopsys Optical Solutions Group | CODE V | LightTools | LucidShape | RSoft Products. 2023. <https://www.synopsys.com/optical-solutions/about.html>.
[24] Photonics Simulation Software | Ansys Lumerical. 2023. <https://www.ansys.com/products/photonics>.
[25] Cadence Photonics: Integrated Design Flows for Photonic Circuits. 2021. <https://www.cadence.com/en_US/home/solutions/photonics.html>.
[26] Photon Design—Your source of photonics CAD tools. <http://www.photond.com/index.htm>.
[27] W. Bogaerts, L. Chrostowski, Silicon photonics circuit design: methods, tools and challenges, Laser Photon. Rev. 12 (4) (2018) 1700237. Available from: https://doi.org/10.1002/lpor.201700237.
[28] B. Waschneck, L.W.F. Brian, K.C.W. Benny, C. Rippler, G. Schmid, Unified frontend and backend Industrie 4.0 roadmap for semiconductor manufacturing, (2017) 20.
[29] JePPIX roadmap 2018, *JePPIX*. <http://www.jeppix.eu/vision/>.
[30] S. Latkowski, D. Pustakhod, M. Chatzimichailidis, W. Yao, X.J.M. Leijtens, Open standards for automation of testing of photonic integrated circuits, IEEE J. Sel. Top. Quantum Electron. 25 (5) (2019) 1–8. Available from: https://doi.org/10.1109/JSTQE.2019.2921401.
[31] M. Trappen, M. Blaicher, P. Dietrich, C. Dankwart, Y. Xu, T. Hoose, et al., 3D-printed optical probes for wafer-level testing of photonic integrated circuits, Opt. Express 28 (2020) 37996–38007.
[32] openEPDA Overview—openEPDATM documentation. 2022. <https://openepda.org/>.
[33] JePPIX Roadmap 2021–2025, *Jeppix*. <https://www.jeppix.eu/knowledge/roadmap/latest-edition/>.
[34] JePPIX Pilot Line—Testing, Jeppix. <https://www.jeppix.eu/jeppix-pilot-line/services/testandassembly/>.
[35] P.I. Dietrich, M. Blaicher, I. Reuter, et al., In situ 3D nanoprinting of free-form coupling elements for hybrid photonic integration, Nat. Photon 12 (2018) 241–247.
[36] E. Perez, G. Moille, X. Lu, D. Westly, K. Srinivasan, Automated on-axis direct laser writing of coupling elements for photonic chips, Opt. Express 28 (2020) 39340–39353.
[37] N. Lindenmann, G. Balthasar, D. Hillerkuss, R. Schmogrow, M. Jordan, J. Leuthold, et al., Photonic wire bonding: a novel concept for chip-scale interconnects, Opt. Express 20 (2012) 17667–17677.
[38] M. Billah, M. Blaicher, T. Hoose, P. Dietrich, P. Marin-Palomo, N. Lindenmann, et al., Hybrid integration of silicon photonics circuits and InP lasers by photonic wire bonding, Optica 5 (2018) 876–883.
[39] M. Blaicher, M.R. Billah, J. Kemal, et al., Hybrid multi-chip assembly of optical communication engines by in situ 3D nano-lithography, Light. Sci. Appl. 9 (2020) 71.

[40] A. Barbot, M. Power, F. Seichepine, G.Z. Yang, Liquid seal for compact micropiston actuation at the capillary tip, Sci. Adv. 6 (2020) 22.
[41] S. Reede, I. Eichhorn, M. Oellers, A. Schander, M.J. Vellekoop, Two-photon polymerized flow sensor integrated in a microfluidic channel with optoelectronic readout, in: 2020 IEEE Sensors, 2020, pp. 1–4.
[42] R.D. Baker, T. Montenegro-Johnson, A.D. Sediako, M.J. Thomson, A. Sen, E. Lauga, et al., Shape-programmed 3D printed swimming microtori for the transport of passive and active agents, Nat. Commun. 10 (2019) 4932.
[43] R. Hensel, K. Moh, E. Arzt, Engineering micropatterned dry adhesives: from contact theory to handling applications, Adv. Funct. Mater. 28 (2018) 1800865.
[44] D. Cheng, R.K. Jayne, A. Tamborini, J. Eyckmans, A.E. White, C.S. Chen, Studies of 3D directed cell migration enabled by direct laser writing of curved wave topography, Biofabrication 11 (2019) 021001.
[45] S. Bertels, M. Jaggy, B. Richter, S. Keppler, K. Weber, E. Genthner, et al., Geometrically defined environments direct cell division rate and subcellular YAP localization in single mouse embryonic stem cells, Sci. Rep. 11 (2021) 9269.
[46] J. Harberts, C. Fendler, J. Teuber, M. Siegmund, A. Silva, N. Rieck, et al., Toward brain-on-a-chip: human induced pluripotent stem cell-derived guided neuronal networks in tailor-made 3D nanoprinted microscaffolds, ACS Nano 14 (10) (2020) 13091–13102.
[47] M. Blaicher, 3D additive microfabrication for relaxed optical packaging (Inv. Paper) Photonics West, 2020 (11286-43).
[48] H. Gehring, M. Blaicher, W. Hartmann, P. Varytis, K. Busch, M. Wegener, et al., Low-loss fiber-to-chip couplers with ultrawide optical bandwidth, APL Photon. 4 (2019) 010801.
[49] H. Gehring, A. Eich, C. Schuck, W. Pernice, Broadband out-of-plane coupling at visible wavelengths, Opt. Lett. 44 (2019) 5089–5092.
[50] H. Gehring, M. Blaicher, T. Grottke, W.H.P. Pernice, Reconfigurable nanophotonic circuitry enabled by direct-laser-writing, IEEE J. Sel. Top. Quantum Electron. 26 (5) (2020) 1–5.
[51] C. Maibohm, O.F. Silvestre, J. Borme, M. Sinou, K. Heggarty, J.B. Nieder, Multi-beam two-photon polymerization for fast large area 3D periodic structure fabrication for bioapplications, Sci. Rep. 10 (2020) 8740.
[52] L. Yang, F. Mayer, U.H.F. Bunz, E. Blasco, M. Wegener, Multi-material multi-photon 3D laser micro- and nanoprinting, Light: Adv. Manuf. 2 (3) (2021) 296–312.

This page intentionally left blank

Reliability of Photonic-Integrated Circuits for data center and high-performance computing applications

13

Gloria E. Hoefler[1], Ranjani C. Muthiah[2], Randal A. Salvatore[1], Peter De Dobbelaere[3] and Robert W. Herrick[4]

[1]Infinera Inc., Sunnyvale, CA, United States, [2]Marvell Technology Inc., Santa Clara, CA, United States, [3]Cisco Systems Inc., Carlsbad, CA, United States, [4]Intel, Santa Clara, CA, United States

13.1 Introduction

Since the 1980s the optoelectronics industry began applying methods from the electronics industry to do reliability testing and prediction [1]. Today there are proven standards, predictive models and relevant tests to address the reliability of optical components and photonic-integrated circuits (PICs) deployed in telecommunication networks [2]. However, the requirements for photonic components and PICs deployed into data center (DC) and high-performance computing (HPC) applications differ significantly from telecom applications. They operate under different environmental conditions, and hence the methodology and standards are adapted to realize and deploy reliable PICs. In this chapter, we draw from the experience of commercial optical networks and describe specific environment and reliability standards and protocols used to deploy reliable PICs.

Similar to the advantages of integration in the microelectronic industry, photonic integration offers advantages in performance, size, power and operational efficiency. Since the early 1990s, the trend in optical communications has been to utilize increased photonic integration in the network to scale bandwidth while simultaneously providing a better system cost than comparable discrete components [3]. The demand for increased transmission capacity paved the way for breakthroughs in photonic integration to overcome transmission impairments due to chromatic and polarization mode dispersion in fiber at high data rates and, later for the deployment of the first fully integrated on−off keying (OOK) transmitters [4,5]. In 2004 we saw the first fully integrated multichannel optical system-on-chip (SOC) 100 Gb/s transmitter and receiver PICs [6]. These early PICs integrated ∼50 distinct functions in one chip. Integration resulted in substantial cost savings and reliability improvements by eliminating fibers within the link as the opto-mechanical

Integrated Photonics for Data Communication Applications. DOI: https://doi.org/10.1016/B978-0-323-91224-2.00010-2
© 2023 Elsevier Inc. All rights reserved.

fiber coupling joint is typically one of the least reliable elements of a photonic component [3,7].

The most widely used integration platforms for optical links are silicon photonics (SiPh) and InP. The salient fabrication and manufacturing processes for both platforms are described in more detail in Chapter 11. Both integration platforms have trade-offs that are addressed at design process to mitigate their impact to reliability and quality. Beyond yields, fabrication process fluctuations and drifts impact the reliability of PICs as they can introduce variation at multiple steps of the manufacturing flow. This chapter will address the on-going process control, monitoring and wafer qualification required to maintain quality.

Unlike discrete components with distinct failure or degradation mechanisms, photonic integration brings about electrical and thermal crosstalk which can impact the system's electrical and optical performance and reliability. For instance, the operating points in multichannel PICs can change as neighboring channels are adjusted or tuned. In addition to the individual component degradation, the influence of crosstalk must be understood to determine the overall performance degradation and lifetime of PICs. Furthermore, in the field, measured changes in performance can result from device degradation as well as servo adjustments. These changes are described in more detail in Chapter 9. One key advantage of photonic ICs is the ability to integrate sensing and control elements to access and monitor variations of key performance parameters during field operation. Such data can be collected for the purposes of monitoring and implementing feedback loops, as well as refining reliability models. The overall benefits of integration, the added functionality and control, typically outweigh any additional complexity.

Another important aspect of integration and reliability is the compatibility of materials used to complete back-end and assembly processes for both InP and SiPh. Whether integration is achieved in SiPh or InP, thermal issues are increasingly important in compact, high-speed optical links, and directly affect the choice of materials. As more and more functions are densely packed, careful consideration of thermal conductivity, the coefficient of thermal expansion (CTE), and general heat management become critical aspects of designing and fabricating reliable PICs. The heat conduction paths in PICs needs to be efficient and stable over the lifetime of the product, and the choice of assembly process and packaging materials must be compatible with all the module components.

To make reliable PICs with high densities of integrated functions and channels, it is essential to use design-for manufacturing (DFM) and quality (DFQ) methodologies with extensive process characterization and quality controls. Each channel in the PIC must meet all requisite system metrics and deliver this performance across multiple channels without incurring a substantial penalty. With the help of DFM and DFQ programs, reliable PICs have been fabricated and deployed in telecommunication networks since the early 2000s. The key metric, of course, are field failure rates. Here we point to commercially deployed 100-Gb/s OOK SOC DWDM PIC pairs (transmitter and receiver) that have clocked more than 3 billion accumulated field hours with zero fails, and a failures-in-time (FIT) rate of <0.3 at 60% confidence level (CL) [8].

Another important data point to attest to the reliability of photonic integration comes from the realization of coherent modulation systems as these architectures employ approximately $4-5\times$ more components compared to the equivalent OOK modules [8]. In 2011 the first 500-Gb/s PM-QPSK transmitter multichannel system on chip (SoC) monolithically integrating $\sim$440 function per chip were deployed in optical networks [8]. These SoC PICs use the integrated control and sensing elements during operation. A constant power output between channels and over time is maintained in the field. Fig. 13.1 shows 44k field hours (>5 years) from a commercially deployed 10-channel 500 Gb/s transmitter module. The attenuation characteristic curve of the integrated VOA located downstream from the Mach–Zehnder modulator (MZM) (see inset), enables the calculation of the change in optical output power from each channel relative to the initial power recorded at deployment to within an Error of Measurement (EoM) of $\sim$0.05 dB, as shown in Fig. 13.1A. The change in transmitter output power shown in Fig. 13.1B demonstrates each of the 10 channels exhibits stable performance, with the worst channel being within 0.4 dB of its initial value [9]. These trends underscore the advantages and capability of integrating sense elements to control and stabilize operation over life for each channel coupled into the fiber.

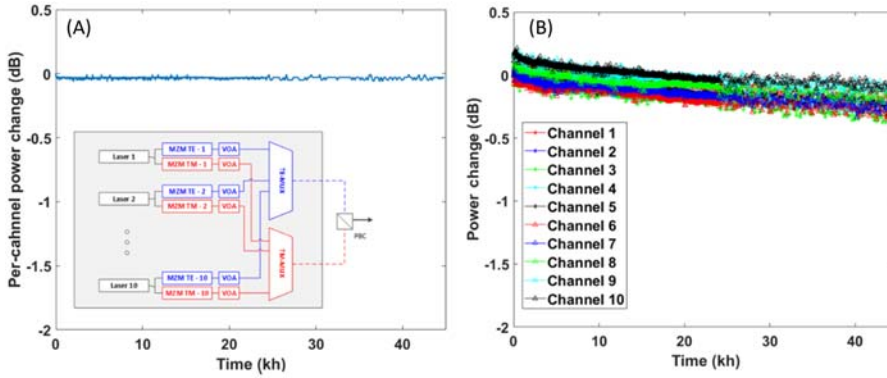

Figure 13.1 (A) Total power stability from Infinera's field deployed 10-channel 500-Gb/s transmitter module measured at the output of the polarization beam combiner. The inset shows the 500-Gb/s PIC architecture; (B) power change on a per channel basis from the same field deployed 10-channel 500-Gb/s PIC transmitter module showing the change in transmitter power measured prior to the VOA.

In the following sections we will describe the methodology used to achieve reliable PICs and provide examples from commercially deployed PICs and modules.

13.2 Reliability requirements for DC/HPC application environment

The reliability requirements for components depend critically on a number of factors. Validating reliability with large qualification experiments and adding screens

to eliminate rarely-seen failures add cost. It is important to understand the applications to know what the tradeoff is between enhanced reliability on one hand and added cost on the other. While it is not a complete list of considerations, some of the more common ones are listed below in Table 13.1.

Table 13.1 Five important drivers of reliability requirements.

Element	Most demanding	More demanding	Less demanding
Business impact of failure	Safety critical links	Online banking	Bitcoin mining; promotional e-mail generation
Are backups communication paths available?	No backup links; single point of failure	Backup links, but require software "repair"	Multiple backups; rapid and seamless switchover in case of failure
Cost of repair	System not repairable	Soldered / embedded parts	Front-panel swappable in small system with embedded diagnostic software.
System size / complexity	Large racks with hundreds of transceivers	Pluggable active optical cables	Single box with front panel swap
Mission Life	20-25 year life for infrastructure	5-7 year for Data Center interconnect	2-5 year for Data Center compute

The drivers of system reliability requirements come from the business impact of the failure. These include items like the loss of business or impact to the customers (on the "lost revenue" side), as well as repair costs on the "increased overhead" side.

The three most important inputs a reliability team needs from the customer are:

1. What is the thermal environment the product is likely to be exposed to?
2. What is the intended deployment life of the product before it is likely to be retired?
3. What failure rate is acceptable to the customer, both in early-life failures, as well as after the first year of deployment.

In this section, we will discuss the answers for typical cloud data center customers.

Thermal environment: Over the past 25 years, cloud data center operators have moved to reduce cooling costs, trying to minimize the consumption of power used for air conditioning. The motivations are both for cost reduction, but also to minimize the environmental impact. Many cloud data centers are placed in naturally cool areas, or near large bodies of water suitable for cooling. Evaporative cooling is also used, and increases the natural humidity in the data center. The ASHRAE and other standards have "opened up the window" to higher temperature and higher ambient humidity than were recommended in the Telcordia GR-63 window [10]. However, the temperature range of ambient air still is much reduced relative to the uncontrolled outdoor exposure (abbreviated "UNC" in the GR-468 specification) that could theoretically be possible. Furthermore, fiber optic transceivers are often "faceplate pluggable," with cool air being drawn in directly over them. In those applications, data on their temperature profile can be directly collected using digital diagnostics monitoring (DDM), which is collected internally, but reflects the

product case temperature. In most cases, this is 30°C−45°C for the case temperature, with a laser chip substrate temperature ranging 10°C−25°C hotter due to self-heating. This is still quite a bit lower than the specified range of deployment temperature that ranges up to 70°C for the case temperature, and 80°C for the chip temperature.

Deployment lifetime: The intended deployment life in the older GR-468 standard was for 20−25 years, which drove the use of conservative packaging technologies, including expensive hermetic packaging, using metal packages, ceramic feedthroughs, and glass windows. Over the past 20 years, there has been a move to less expensive nonhermetic packaging, using primarily plastic, or unenclosed components. These have some risk of corrosion-based failure mechanisms, particularly in long deployments, as well as worst-case outdoor deployments. An example would be a 5-G data terminal deployed on a mast in a tropical location, where the intended service life might be up to 15 years. But for cloud data providers, the guidance is that they normally plan on a 5-year deployment lifetime, and occasionally extend to 6 or 7 years, but never beyond that duration. They are more interested in potential cost reductions from nonhermetic packaging than they are in higher cost designs that support 10 year or longer deployment lifetimes.

Acceptable failure rate: Acceptable failure rates have risen somewhat over the past 20 years on a per module basis, although they have dropped rapidly on a per-GB basis. For the current generations of 100−800 Gbps transceivers, the large cloud providers are citing goals of 300 FITs per transceiver (roughly 0.3% per year failure rate). In addition, the "dead on arrival" failure rate is specified as being under 500 defective parts-per-million (dppm), or 1/20th of a percent [11].

13.3 Reliability standards for DC and HPC applications

Optoelectronic device reliability is guided by the Telcordia Core GR-468 rev 2 standard [2]. The standard provides a number of required reliability tests which must be completed on the transceivers and documented in a reliability qualification report. The tests can be broken down into four categories:

1. **Endurance tests**: these are ones at elevated temperature that accelerate the degradation of the parts. They include things like "high temperature operating life" (HTOL), and "high temperature storage" (HTS).
2. **Corrosion tests**: these involve accelerating corrosion failure mechanisms by aging parts, either with or without bias, in environmental chambers with high temperature and high humidity. The most common test will be the "85/85" test, which is done in a chamber at 85°C, with 85% relative humidity. For most failure mechanisms, this accelerates corrosion by approximately 100× over worst case tropical outdoor deployment. In larger test matrices, these could also include exposure to corrosive gases.
3. **Mechanical tests**: these include shock and vibration tests to verify that normal shipping won't damage the parts. They also include "temperature cycling" tests, where the parts are alternately cycled between −40°C and 85°C hundreds of times. Due to the different components in the transceiver having different coefficient of thermal expansion, these

tests address whether the design is vulnerable to fatigue cycling from the mechanical stress seen due to the mismatches between the coefficients of thermal expansion. Mechanical fatigue can occur from the cycling of the stress, often resulting in solder cracking or component failure in cases where the device failure is seen.

4. **Special tests**: the standard requires electro-static discharge (ESD) testing, to establish how robust the transceiver or components are against damage from handling. Components that are very sensitive would either need to be re-designed, or special handling precautions would need to be taken if this was not possible.

When a new component is being designed, early prequalification testing is ideally done as soon as prototypes are available, to expose any design weaknesses. The sample size for each test may vary between 11 and 25 devices. However, the formal qualification matrix following JEDEC/GR-468 protocols normally requires at least 77 devices, to do most legs in parallel, rather than sequential tests, and would require 5−9 months, depending on whether the less demanding 2000-hour duration, or the more demanding 5000-hour duration version of the qualification is being sought. (The former would be for "controlled operation"—environmentally controlled indoor only, while the latter would be for "uncontrolled operation," with remote terminals that are not environmentally controlled.)

Higher level assemblies (such as modules) use additional testing, including blowing dust tests, fiber integrity tests which pull or bend the fiber connectors, and insertion-removal tests for fibers and electrical connectors. They also may do other tests, such as thermal shock, and cyclic moisture resistance, where the latter absorbs moisture during a "soak," and then is followed by a "freeze," along with power-up cycles.

13.4 Methodology for PIC reliability

In Fig. 13.2, we outline the key steps for a robust design for quality program (DFQ).

Chapter 10 describes the approach and tools used to implement a design for manufacturing (DFM) for photonic ICs. The subsequent sections will cover details on the key steps of the DFQ program and the interaction with the design process.

13.4.1 Failure mode effects analysis

The failure mode effects analysis (FMEA) is an integral part of the New Product Introduction (NPI) program used to assess and quantify reliability risk in relation to the field operating conditions. The FMEA methodology is performed to systematically identify the ways in which a product or process can fail, estimate the risk associated with failures, and prioritize the actions that should be taken to reduce risk. FMEA are conducted as a cross-functional iterative activity with all key stakeholders of the product. The risk for a given failure mode are captured in the risk priority number (RPN) which is calculated as the product of three aspects—the severity (or field impact), the occurrence rate, and the detection mechanisms that are in-place for that failure mode.

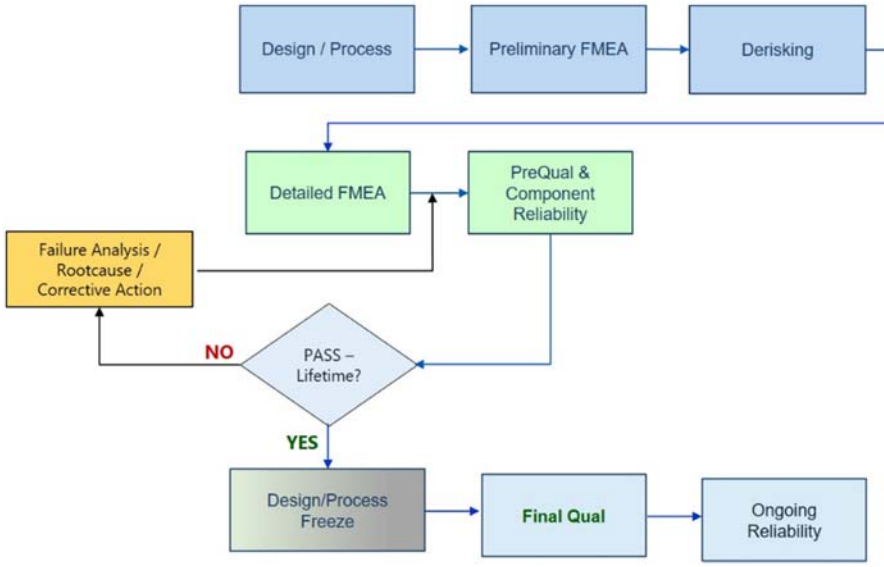

Figure 13.2 Key steps for a design for quality programs.

There are two types of FMEAs performed: (1) design FMEA (DFMEA) for design-related failure modes performed early in the NPI Phase and (2) process or manufacturing FMEA (PFMEA/MFMEA) for process related failure modes performed early in the NPI Phase and closer to the production release of the part. The output of the FMEA is a prioritized list, based on RPN determination, of derisking and prequalification activities conducted to assess and enhance the quality and reliability of the product. The reader is referred to [12] for a comprehensive review on the FMEA approach and reference [13,14] for standards for the microelectronics.

13.4.2 Derisking

The derisking plan, also called STRIFE test (stress and lifetime test), is an essential activity in the early phase of the DFQ program for all key elements prior to an integrated PIC qualification. The goal of derisking tests is to identify/understand and eliminate key failure modes and acceleration factors for each element of the PIC, critical component, and subassembly of the module. The results from these tests are used as inputs into the final FMEA and help define the prequalification/qualification and component reliability test plans. Any failures resulting at this stage are analyzed and the learnings fed back to the development teams to improve the design and process for the product.

Derisking tests are performed at various levels —individual elements level where test structure of a PIC are isolated, integrated PIC level where individual or multiple elements are stressed, subassembly level where multiple components are integrated along with the PIC on a substrate or test package, and early module level where the entire assembly can be stressed.

Quantifying the level of stress exposure any element type (laser, tuning element, detector, modulator, splitter, combiner, waveguide, facet, etc.) can sustain without degradation in function is critical to obtain relevant lifetime estimates. Multiple tradeoffs between the duration of the test, cost, time and stress conditions are considered to avoid irrelevant failure modes when deciding the conditions for reliability testing. Degradation can be accelerated in many dimensions—by elevating temperature, increasing current, increasing voltage, increasing humidity, increasing thermal gradients, increasing optical power, through thermal cycling, through mechanical shock, through thermal shock, and more [15]. Some degradation effects have gradual, continuous trends that can be measured and tracked, while others are sudden, holding stably to the expected operating point until they abruptly fail. Continuous effects can be more readily predicted with smaller sample sizes and longer-spaced sampling intervals, while the sudden failures only provide binomial statistics which require large sample sizes in order to achieve significant confidence levels [16]. In cases where there is ample margin between required field operating condition and the point of noticeable degradation, the component reliability can be met by design. In practice, to achieve maximum field performance, most system and device designers drive to operating conditions (temperature, voltage, current density, etc.) near the level where degradation begins to occur.

When introducing a new product or component an initial step-stress is performed to determine the safe bias conditions for the stress testing. To illustrate this point Fig. 13.3A shows a high-level view of a step-stress testing flow for a photodiode (PD). The dark current (Idk) versus step voltage (V) shown in Fig. 13.4B is used to determine if there is sufficient margin between the PD's operating voltage (Vop) and its breakdown voltage. The dark current at the use voltage is shown for reference (red line), the data indicates it is safe to bias to 2*Vop, and even 3*Vop, for stress testing. This graph also illustrates the onset of significant PD leakage at higher voltages (>2*Vop), this region can be used to detect unacceptable process changes in the PD.

To quantify the PIC stability over life, it is necessary to use accelerated temperatures that do not result in additional failure modes yet provide small degradation changes. This approach leads to increased STRIFE test time to obtain adequate signal to noise for the degradation. Ultimately, to get acceleration factors, elevated currents and temperatures are required. Typically, changes in solids that are driven by microscopic events, such as defect propagation or chemical reactions, can be characterized by a thermal activation energy. However, some types of degradation related to macroscopic changes in material property, phase transitions, etc. have characteristic failure threshold temperatures and cannot be represented by a simple activation energy. Accelerated tests should be conducted with sample sizes large enough to have confidence in the lifetime predictions and the resulting activation energies. From an end-user perspective, it is not sufficient to review a vendor's minimal qualification data, revealing only standardized pass/fail statistics for a single time interval as these simple reports give low confidence regarding the failure mode and mechanism. Instead, an end user prefers to see challenging STRIFE tests that have driven components well into their failure regime with extensive exposures

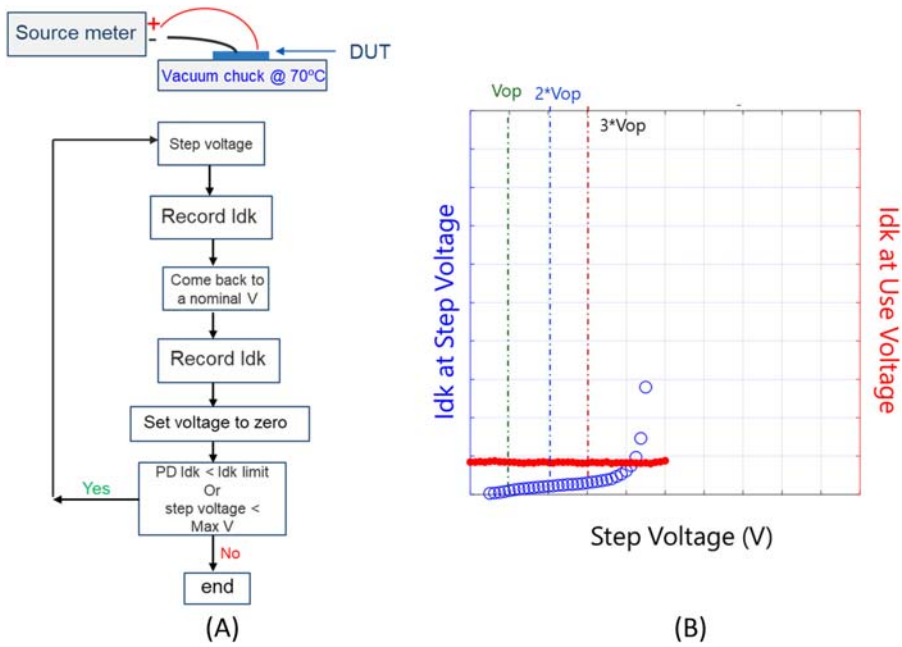

Figure 13.3 (A) Step stress flow for a photodiode (*PD*); (B) dark current (Idk) as a function of step voltage (*V*) the measured Idk at use voltage is shown for a Silicon photonics PD.

that reveal how components are expected to fail along with a sound model of the failure limits, the degradation trends, and the root cause.

However, excessive acceleration conditions can obscure the ability to predict the failure rate and mode that is relevant to the field operating conditions [17]. Both the dominant sudden failure mode and the dominant wear-out failure mode need to be quantified while keeping in mind that a "dominant failure mode" is specific to the field operating conditions and that another failure mode can dominate at highly accelerated conditions. Failure analysis (FA) can be used at the derisking phase to assist in determining the root cause of the failure mechanism (mechanical, physical, chemical or other) and assign relevance when there is more than one failure mode accelerated. The FA process is multidisciplinary in nature and is performed on all failed parts. FA can include destructive techniques nondestructive techniques, such as visual inspection, optical microscopy, X-ray radiography, chromatography, simulation tests and/or decapsulation, sectioning, microprobing, and chemical analysis.

Due to the nonlinear behavior that results from high acceleration of components with different materials, care must be taken to confirm the trends obtained with accelerated stress are extrapolated to the representative field condition for the product. One approach to mitigate the risk to lifetime predictions is to run tests at different temperatures for the same total time. Fig. 13.4 illustrates a study where a PIC parameter was characterized after 2000 hours of accelerated stress. Under condition "Temp4" a very large decrease in this parameter was seen while a lower condition, "Temp3", showed

much smaller decrease over time. Applying an Arrhenius fit to Temp3 and Temp4 conditions alone, the acceleration factor would project a high activation energy and a very small change for the field condition, as shown by the dash-dot green line in Fig. 13.4. Yet, data from two, lower temperature conditions, Temp1 and Temp2 resulted in different sign of degradation. This indicated that conditions Temp3 and Temp4 were well beyond the linear degradation level. The blue solid line shows a fit to the experimental data requiring a crossover temperature Tco (temperature at which two processes equally contribute) with the high temperature degradation process having a considerably higher activation energy than the low temperature degradation process.

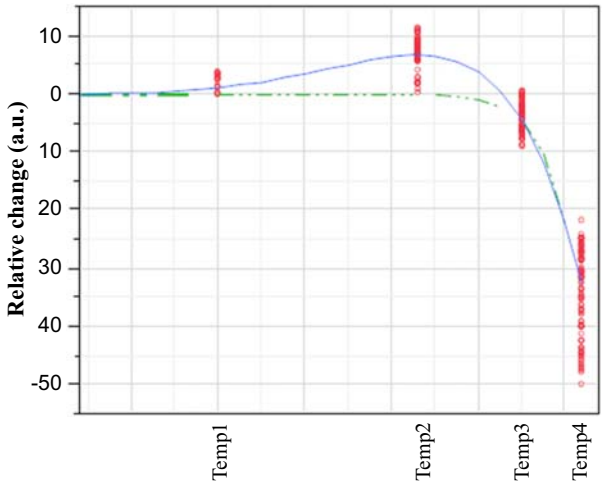

Figure 13.4 Change through 2000 h of aging versus temperature for four different temperature conditions (Temp1−Temp4). The red circles correspond to aging data, the green dash-dotted line is an Arrhenius fit based on the upper two temperatures (Temp3 and Temp4). The blue line corresponds to a fit based on (T-Tco) × Arrhenius acceleration, where there is a change in the impact from a different degradation process stress at characteristic temperature Tco.

It is noteworthy that all of these accelerated conditions (Temp1, Temp2, Temp3, and Temp4) gave continuous degradation trends at a level below any indication of sudden degradation failure modes. Using the most aggressive condition alone would have predicted not only the wrong magnitude but also the wrong direction, less than 0.1 units of degradation in the field over 7 years resulting in overly optimistic reliability projections. This example demonstrates the importance of mapping out the trends to the field conditions when using highly accelerated cells.

13.4.3 Design constraints

Because overall system-level failure rate is the aggregate of the failure rates of all the individual components of the system, each element impacts the net failure rate. To some extent, integration enables a trade-off between the failure budget of individual components to maximize the overall system performance. The caveat here, is

that any component with higher failure rate or variation requires more stringent monitoring. Potentially higher variation in failure rate is especially problematic in applications of DCI and HPC where rapid deliveries and large volumes would necessitate a robust quality assurance program to stay within a specified FIT budget.

The commonly used mean time between failure (MTBF) is not a meaningful metric for optical communication networks because the distribution of failures in time for all the components must also be understood. One may project MTBF of 100 years for a component, but a 5% fail over the first year on any component can have a big impact on the integrity of the network. Instead, the focus is on the failure rate as a function of time and specifically on the single digit percentage region of the failure curve. Hence a more appropriate metric is the failure in time (FIT) distribution of components, where FIT rate is the fraction of fails expected for 1 billion devices-hours, or alternately, is a parts per million failure rate per kilohour of aging time. Note that the FIT rate (or Hazard rate), may vary over time [2].

A DFQ program requires designers to budget for a FIT rate that can be maintained over the product lifetime. However, at the time of product release, even with extensive reliability characterization and accelerated aging, due to reliability test sample-size limits, it is typical to only have data to certify that the failure rate is less than 500−1000 FIT, and to require additional monitoring during the initial product shipment ramps. Practical knowledge of a transceiver with $<$100 FIT only comes from field data after a successful, long-term field deployment. Thus in practice most successful new products with low FIT are based on extensive reliability models based on prior generations with upgrades to some of the components.

Although DCI and HPC applications have a shorter lifetime requirement that mitigates wear-out failures, it however does not in any way reduce the requirements for infant mortality or random failures. Overall, the reduction from a $\sim$15-year lifetime requirement to a $\sim$7-year lifetime requirement, minimally reduces the qualification effort needed to achieve low FIT rate over the 7-year life.

In summary, the qualification team needs to map failure modes as a function of use environment for all critical components. As Fig. 13.2 indicates, a DFQ program requires multiple iterations on the design, process and operating requirements to meet lifetime requirements. Not only must the specifications be met for a certain era of PIC, but the quality of these components must stay consistent, and the acceleration factors for each component must remain applicable throughout the manufacturing period of the product. This requires the original derisking and qualification data to be based on a diverse sample from multiple fab lots, multiple wafers and regions on these wafers and for key process monitors and controls to be in place to assure the process does not deviate from the qualified process.

13.5 Subsystem and module assembly and integration

Hyperscale data center optical interconnect requirements, such as data rate, link performance, power dissipation, and density, are currently driving advanced optical

transceiver technology. Highly integrated optics solutions enable this trend. In this section, we will discuss reliability aspects of optical engines. We will address the reliable integration of photonics and electronics, light sources in silicon photonics and the coupling between the PIC and single mode optical fibers for both SiPh and InP integration platforms.

13.5.1 Reliability of photonics/electronics integration

In optical transceivers the combination of optical and electronic functionality is essential. Some early silicon photonics implementations had photonics (e.g., modulation, photodetection) and electronic functions (modulator driver, trans-impedance amplifier) integrated in a single IC [18,19]. In this case, reliability of the integrated device was ensured by adhering to the reliability design rules (e.g., electromigration, metal seal ring at the edge of the die) of the IC technology that was used to manufacture the electro-optic IC. As data rates increased, electrical and optical functionalities were split out between and electronic IC (EIC) and a PIC to save product and development cost. Current EICs are made in advanced CMOS nodes and besides modulator drivers and trans-impedance amplifiers also have additional functions, such as SerDes, digital signal processing, and controls. The interconnection between EIC and PIC is a key contributor to the performance of the chipset. Flip-chip bonding by means of microbumps, such as Cu Pillar technology, shown in Fig. 13.5A is now frequently used to enable low parasitic (R, L, C) high density electrical interconnect. This mature technology is leveraged from the electronic IC industry. Outsourced Semiconductor Assembly and Test (OSAT) providers have established design rules for high-yield and high-reliability.

Over the years a number of reliability concerns related with microbumps flip-chip bonding have been addressed:

1. Under-bump metallization, diffusion barriers and solder compositions have been optimized to prevent formation of intermetallic compounds and resulting Kirkendahl voids [20,21] which can mechanically and electrically weaken the bond reliability.

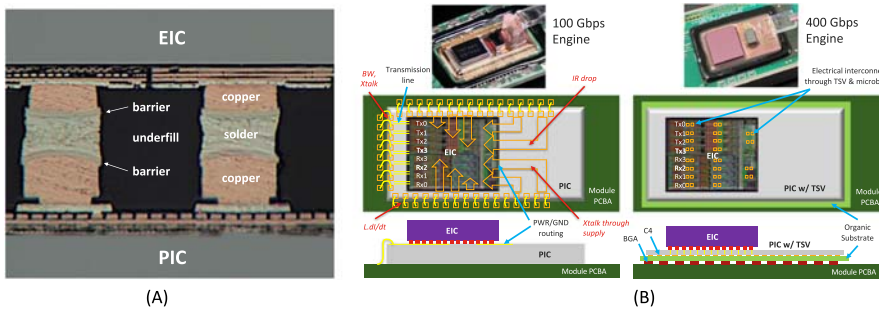

Figure 13.5 (A) Microbump interconnection between EIC and PIC (25-μm diameter, 50-μm pitch). (B) Wire bond versus integration by 3D integration by TSV. Wire bonds result in inductive, capacitive, and resistive parasitics that penalize signal integrity and power.

2. Materials have been selected to minimize the effects of differential thermal expansion coefficients between the bonded ICs and substrates. In the case of silicon photonics engines, this is less of a concern as both substrates are Si.
3. Materials and design rules have been established to minimize electromigration in the interconnect [22]. The importance of this mechanism might increase as more advanced CMOS nodes evolve towards lower supply voltages and hence higher supply currents.
4. Suitable underfill resins are applied between the two ICs to provide additional mechanical stability and protection.

Earlier implementations of microbumps allow wire bonding from the PIC, which is usually larger, to the substrate. However, as density and complexity of transceiver functions increase, wire bonding becomes less suitable because of the inherent parasitics of the bond wires and corresponding pads and electrical routing on the PIC affecting both signal integrity for the high-speed signals as well as power supply. Therefore new technologies such as 3D integration are employed to overcome those issues. Fig. 13.5B shows how implementation of Through Silicon Vias (TSV) realizes low parasitic interconnect directly through the silicon PIC substrate to microbumps connecting specific nodes in the EIC by eliminating wire bonds.

This approach leverages from the IC industry where TSVs are currently used in silicon interposer fan-out assemblies [23]. Intrinsic reliability concerns within in the TSV structure itself (e.g., electromigration and dielectric breakdown) and chipset assembly have been addressed and suitable standardized qualification requirements have been established to ensure reliability [24]. Si foundries and OSATs have established design rules for design and manufacturing of 3D chipsets [25]. Fig. 13.6 shows a typical chip-on-chip-on-substrate (CoCoS) photonic assembly consisting of: (1) a thinned down PIC ($\sim$100 μm) provided with TSVs bonded to an organic substrate by C4 bumps on which (2) an EIC is flip-chip bonded by microbumps. The warpage of the PIC and PIC-on-Substrate during assembly soldering may cause reliability concerns. Therefore strategies based on stress compensating layers have been developed to keep this warpage under control and ensure reliable bonds within

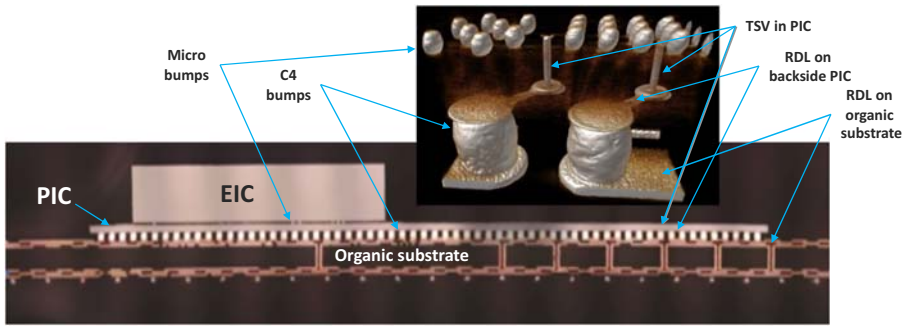

Figure 13.6 Cross-section of 3D-integrated silicon photonics CoCoS assembly. Inset 3D X-ray image of interconnects: C4 bump between organic substrate and PIC, TSV through PIC, microbumps between PIC and EIC.

the assembly. An underfill resin is applied respectively between EIC and PIC and PIC and substrate to protect the assembly and bonds [26].

The influence of stress/strain induced by the TSVs on nearby active and passive silicon photonic devices has been investigated to confirm there is no meaningful effect on the performance or reliability of photonic devices integrated on the PIC.

13.5.2 Burn-in

Burn-in is the common practice of exercising a device at an elevated condition (high current and temperature) for a predetermined period of time beyond premature failure, to stabilize and detect early failures prior to assembly or field deployment. Components that passed the burn-in test are more likely to perform reliably over their predicted lifetime. Burn-in can be carried out on lasers, modulators, detectors, or any component to reduce the potential for defects and failures during normal operating conditions. Optical communication laser diodes are typically burned-in by operating for tens of hours in an oven tracking any changes in voltage, frequency, output power before, during, and after the burn-in.

13.5.3 Reliability of light source (not integrated in the PIC)

It is broadly accepted that packaging/bonding of laser diodes is a potential contributor to the failure rate of a laser diode [27,28]. Reliability can be affected by mechanical stress due to the bonding process on the substrate (e.g., AlN or Si) causing migration of crystal dislocations. Therefore it is an industry practice to screen laser diodes for reliability (burn-in) only after they are bonded to their substrate/submount.

Facets of laser diodes are sensitive to environmental conditions, such as humidity. This can be mitigated by using appropriate facet coatings or packaging the laser diode in a hermetic housing. When laser diodes are hermetically packaged, the reliability of the hermetic housing must be verified and the gaseous content inside the package must be free of humidity or other contaminants that may affect reliability of the laser diode. This can be verified by residual gas analysis [29]. An example of a laser diode package suitable for integration with silicon photonics is shown in Fig. 13.7. The laser diodes are mounted to a Si-wafer using pick and place tools. This type of miniature laser diode package is suited for silicon photonics implementations where the optical interfaces to the PIC are based on grating couplers.

The laser micropackage (LaMP) integrates a conventional DFB laser diode, a coupling lens, an optical isolator and a beam steering mirror in a hermetic silicon package that can be manufactured and tested at wafer scale using commercial pick-and-place tools. A 6″ wafer contains up to 3200 light sources. The individual cells are tested for performance (LIV, SMSR, etc.) and screened for reliability (burn-in) by fully automated wafer tools. All 3200 cells are burned-in at the same time. Before and after the burn-in stress key parameters, such as threshold current and

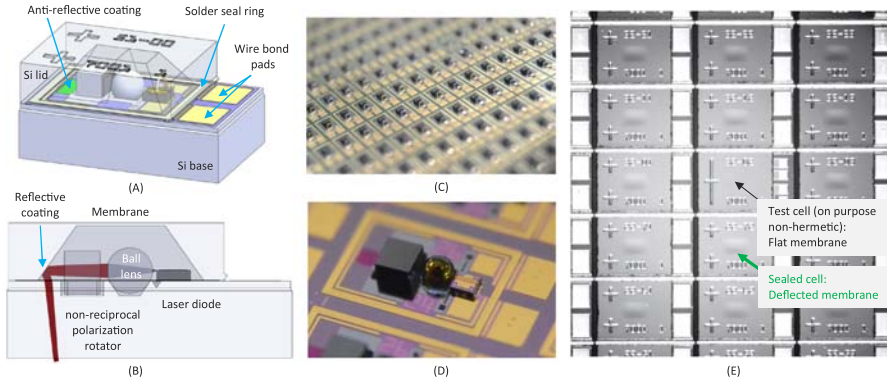

Figure 13.7 Laser micropackage (A) LaMP, (B) optical path in LaMP, (C) micrograph of wafer level assembly of LaMP, (D) higher magnification showing single LaMP, and (E) Vision-based hermeticity test of LaMP.

output power of the laser diodes, are measured using commercial optical wafer probers. If the pre−post change for a cell exceeds a given limit, the cell is rejected. As we are using laser diodes from multiple suppliers the fall-out at burn-in may vary from a few tenths of a percent to a few percent.

Hermetic sealing of each of the cells is obtained by wafer bonding under vacuum using a solder seal. Conventional He bombing/leak testing cannot be used as the inside volume of the LaMP is too small, hence we developed an optical method to confirm hermeticity. The method is based on deflection of the membrane on top of the LaMP due the pressure differential between the inside and outside of the LaMP after sealing. The deflection is a measure for the hermeticity

Ongoing reliability testing includes sample HTOL testing and sample residual gas analysis (RGA) to confirm low moisture content. Our criterion is ≤ 3 molecular H2O layers on the inside surface of the micropackage [30].

13.5.4 Reliability of lasers integrated into the PICs through heterogeneous integration

Lasers integrated into the Si PIC were first demonstrated in 2006 [31], using InP gain epitaxy bonded to a companion silicon photonics wafer, and processed into integrated lasers that are embedded in the platform [32]. They were originally referred to as "hybrid lasers," and we will use that terminology in this section, although current terminology also refers to the "heterogenous integration of indium phosphide with silicon" or "heterogeneously integrated lasers." The technology was commercialized starting in 2017 by Intel [33]. They have subsequently captured a significant market share in the single-mode intra-datacenter transceiver market, since they have been cost-effective, and highly reliable alternatives to transceivers using conventional photonics. Several million transceivers

have been shipped over the past four years, and the resulting Tx PIC failure rate in the fielded application has remained under 2 FITs [33], a figure far lower than that obtained by competing direct modulated laser (DML) technologies. It appears there are principally two reasons for this lower failure rate. First, since modulation is performed separately from light generation, the laser does not need to run at high current densities to get rapid fall times, the way a DML would. Thus current densities are almost an order of magnitude lower in the hybrid laser than in typical DM lasers. Second, much of the failure in DML-based transceivers, which are said to average 300 FIT failure rates (i.e., nearly a 0.3% annual failure rate) with 80%−85% of the failures occurring due to laser failure, comes from climb dislocations that originate at the facet. Cleaving steps are ready sources of micro-cracks, and highly optimized cleaving and coating processes are critical in such lasers. In lasers used for high-power applications, a "window" is formed by regrowth or intermixing to eliminate the active region and current flow for the last few microns before the cleaved facet. However, this adds substantially to the fabrication complexity, reduces the yield, and increases the cost, so it has not been used in transceivers deployed in data center applications. By contrast, the hybrid lasers use a taper at each end of the InP mesa to gradually push the light down into silicon waveguides, so no facet cleaving or coating is involved in this technology, and the nucleation point for defect formation can also be eliminated.

Long-term aging studies have been performed, both at near use condition for a few years, and also at highly accelerated conditions the equivalent of >15 years of normal use, as shown respectively in Figs. 13.8 and 13.9. The InAlGaAs material used in these lasers is known for being very stable against long-term aging.

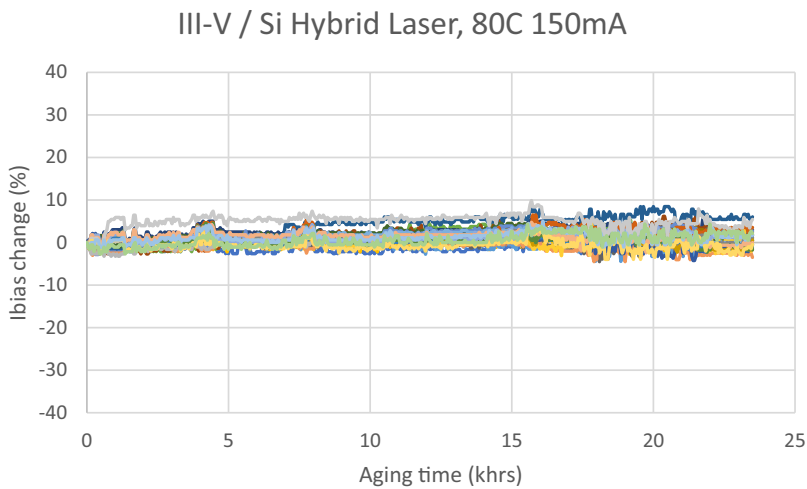

Figure 13.8 Aging on 22 lasers over 23,500 h (2.7 years), at 80 C, and double the normal use current, and shows <10% change in lasing parameters.

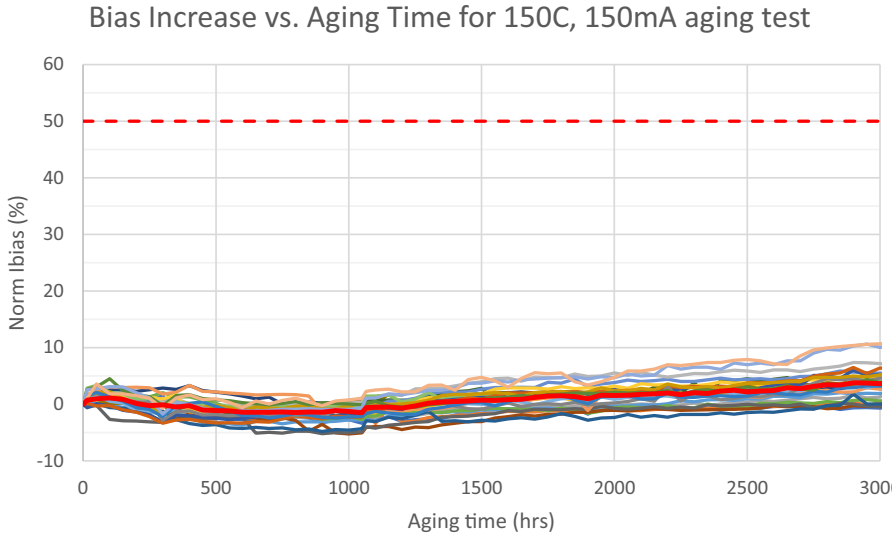

Figure 13.9 High temperature accelerated aging test. Three hundred fifty hours is the equivalent of 10 years at 55°C. Two thousand hours is the equivalent of 10 years at 80°C.

13.5.5 Reliability of fiber interfaces to PICs

Providing one or more single mode optical fiber interfaces to a PIC is a key manufacturing step of integrated optics products. Several technologies have been developed ranging from laser welding to adhesive bonding [34].

Key for the fiber attach is that this is accomplished with the lowest additional insertion loss and that this insertion loss is stable of overall operating conditions and over the useful life of the product. A first step to minimize insertion loss is to match the optical modes between the PIC and the fiber. This can be done in multiple ways at either the PIC side (mode expanders) or the fiber side (lensed fibers). Alternatively, a lens can be used to match the mode sizes, but this results in additional alignments and sensitivities.

It is advised to keep mode sizes sufficiently large ($\sim$9 micron diameter) as for small mode sizes the x-y alignment tolerances may be very tight and become mechanically hard to maintain. On the other hand, if the mode sizes/beam waists get very large (e.g., 100-μm diameter), the angular alignment tolerances get very tight. The reader is referred to Chapter 12 for a more detailed discussion of the tolerances and trade-offs between mode-size and alignment.

Next, we will focus on adhesive bonding of SMF fibers to mode matched PIC interfaces (e.g., grating coupler or edge coupler). Fiber array units (FAU) can be used to conveniently couple multiple fibers to a PIC. In a FAU the fibers are arranged in accurate positions ($<$1 micron) in a glass or silicon block. This is obtained by placing the fibers in precision micromachining grooves in a glass or silicon substrate and fixing them in place with an adhesive. The end face of the fibers

is subsequentially polished to create an optical surface. Such an FAU can then be aligned and bonded to a PIC with an adhesive. Optical bonds can fail catastrophically or gradually. Catastrophic failure may be caused by delamination resulting in mechanical failure, sharply increased coupling loss or back-reflection at the interface. Gradual failure is usually caused by slow increase of coupling loss due to relative movement of the fiber and the PIC. This movement can be caused by, for example, absorption of moisture or thermally driven viscoelastic effects in the adhesive.

Multiple factors may affect the reliability of a fiber to PIC adhesive bond, including:

1. Bonding strength of the adhesive: bond strength comes from either physical or chemical bonding of the adhesive to a surface. Some adhesive formulations include materials, such as organosilanes, that promote chemical adhesion to specific surfaces, such as glass or silicon oxide [35].
2. Cleanliness of bonding surfaces: besides visible contamination, invisible contaminants (e.g., from outgassing of containers) on the surface can negatively affect the strength and reliability of adhesive bonds as they can inhibit chemical bonding of the adhesive. Therefore it is advised to clean surfaces with UV-Ozone or plasma cleaning prior to bonding.
3. Curing condition: Most adhesives can be cured by heat or by light. UV-curing adhesives are particularly attractive in optical assembly as they allow "freezing" an aligned assembly (e.g., FAU and PIC) in place by exposure to UV light. The latter is enabled by adding photo-initiators to the adhesive formulation. Those initiators start the cross-linking process that hardens the adhesive. Several UV curing adhesives exhibit an exothermal reaction during the curing process which may affect alignment during curing. Suitable design and process controls are required to obtain reliable bonds [36,37].
4. Thermal and viscoeleastic properties of adhesives: Cured adhesives are crosslinked polymers and exhibit a glass transition temperature (Tg). At this temperature or temperature range, the polymer goes from the "glassy" state to the "rubber" state and mechanical properties, such as Young's modulus, changes over several orders of magnitude and the thermal expansion coefficient with a factor two or three. Therefore where tight mechanical stability needs to be maintained, one needs to avoid that the glass transition temperature of the cured adhesive is within or near the operating range of the adhesive bond [33]. Tg can be measured by, for example, Differential Scanning Calorimetry (DSC) or Thermo Dynamic Analysis (TGA). The latter technique measures the Young's modulus as a function of temperature. It should be noted that even below Tg there can be viscoelastic phenomena, such as polymer aging, that may affect mechanical properties of adhesives. Those may cause slow degradation of coupling efficiency.
5. Bond line thickness control: Once settled at a given adhesive bond line thickness it is best to keep it within tight control to avoid reliability excursions. Usually very thick or wedged bond lines are to be avoided as they can amplify the effects of viscoelasticity or swelling due to moisture absorption.
6. Thermal and photochemical degradation: Like most polymers, adhesives start to degrade (decompose) at elevated temperatures. Structural adhesives need to be selected so that no decomposition occurs at process and operating temperatures. Adhesives can be evaluated for thermal degradation by thermogravimetric analysis (TGA). Another degradation effect is photochemical degradation where light with a given wavelength (usually the absorption

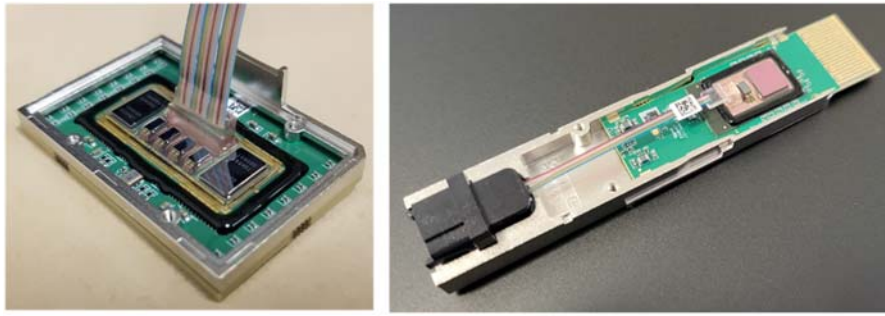

Figure 13.10 Examples of fiber interfaces to silicon photonics PICs. (A) 400 G on-board-optics module (16 × 25 Gbps) with 32 fiber interfacing to the PIC. (B) 400 G Dr-4 QSFP-DD module (4 × 100 Gbps).

peak of the material) affects the chemical composition of the adhesive. Therefore adhesives with absorption peaks (or harmonics thereof) close to the operational wavelength should be avoided.

Fig. 13.10 shows two examples of fiber interfaces as realized in silicon photonic products. In both cases, multiple fibers are accurately kept in place in a FAU bonded to a PIC by means of UV curing adhesive. Optical coupling with fibers occurs through grating couplers realized in the silicon photonics PIC.

13.5.6 Reliability of packaging for photonic subsystems

Previous packaging examples, as shown in Fig. 13.10, demonstrate chip-on-board packaging of silicon photonics subsystems for low-cost, high-volume intra-datacenter interconnect. Such packages have proven to be reliable in the field. Advanced InP PICs are packaged using similar technologies as described earlier in the chapter. However, the monolithic integration of lasers and other gain components enables the implementation of optical links with all the requisite components to transmit and detect while eliminating multiple opto-mechanical fiber coupling joints which are one of the least reliable elements in a photonic package. Infinera reported up to $>100\times$ reduction in fiber couplings in 1.2 Tb/s coherent transmitter and receiver realized with InP PICs [8].

As more and more functions are densely packed, careful consideration of thermal conductivity, the coefficient of thermal expansion (CTE), and general heat management become critical aspects of packaging reliable PICs. The heat conduction paths need to be efficient and stable over the lifetime of the product, and the choice of assembly process and packaging materials must be compatible with all the module components. Infinera's InP transmitter (Tx) and receiver (Rx) PICs and the necessary drive electronics are copackaged in a hermetic high-temperature cofired ceramic (HTCC) land grid array (LGA) package. Hermetic fiber feedthroughs are used to ensure protection of the optics inside the package. Reliable modules are

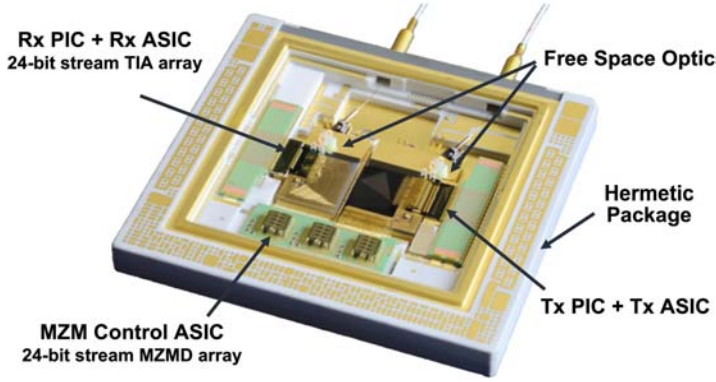

Figure 13.11 Picture of 1.2-Tb/s coherent transceiver module (6 channels x 200Gb/s per channel). The transceiver integrates 6-channel Rx and Tx PICs, a 24-bit stream MZMD and TIA array, a compact multi-chip-module with control ASIC drivers, thermoelectric cooler, free-space optics and two fiber couplings.

achieved by the codesign and thermal management of all the integrated technologies: InP PIC, optical interconnects, high-speed electrical interconnects, DC control and bias circuity. Fig. 13.11 shows a field deployed 1.2-Tb/s coherent transceiver module (lid removed). The transmitter (receiver) photonic ICs are assembled with MZM driver (TIA) array ASICs integrating 24 high-speed bit streams in SiGe BiCMOS technology, thermo-electric (TE) coolers, and DC-stacked control ASICs drivers. The transceiver integrates free-space optics for polarization multiplexing/demultiplexing of the polarization multiplexed signals [38].

The next sections will illustrate the reliable performance of silicon photonic and InP PICs and PIC-based modules in the field.

13.6 Reliability of DCI optical systems and field data

Ultimately, the field data is the true indication of the stability of optical transceivers. Marvell's Silicon Photonics PICs have been deployed in the field since 2016. Since then, multiple generations of PICs have been through successful accelerated reliability testing and extended reliability testing (generally to 5000 hours and sometimes to 10,000 hours). As part of the reliability testing, integrated components are biased at their maximum operating conditions and a sampling of photodiodes are biased at 2 × of maximum V_{op}. The plots in Fig. 13.12 show the electrical characteristics of the photodiode (PD), Mach Zender Modulators (MZMs) and the phase control elements through a 10,000 hours accelerated aging cell. The PIC components are stable through the course of the accelerated aging, which translates to 100 years at an activation energy of 0.7 eV or 15 years at a conservative activation energy of 0.4 eV. There was no difference in the aging of the PDs biased at max. Vop and 2 × of max. Vop. For some types of aging mechanisms, the margin of

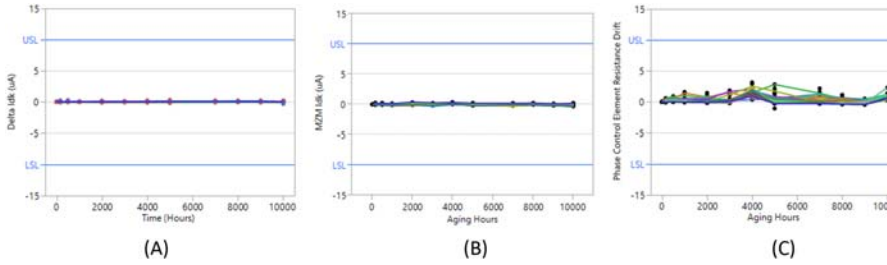

Figure 13.12 Change in dark current (Idk) of (A) PDs, (B) MZMs, and (C) phase control element resistance during 150°C/10,000 h accelerated aging test for Marvell Silicon Photonics PIC.

strength is so large, that it is not possible to collect acceleration factors, such as thermal activation energies. But so long as the projections show an acceptable lifetime even using conservative acceleration factors, it is not required by the GR-468 standard to determine activation energy. The reliability of the PIC was demonstrated by extending the test to a duration equivalent to the use lifetime at conservative activation energies.

Infinera's longest PIC field data comes from the PICs which are 100 Gbs (10 Gbs × 10 channel) optical transmitter and receiver modules, which extended for 133,320 hours, or 15.2 years [8]. During operation, all critical parameters are tracked, as described in the section below. As part of the DFQ process, understanding is gained for how much any component can vary over life, and mathematical models are built, thus sufficient margin can be set aside at beginning of life for any parameter that would need servo control over life. The use of integrated sense and control elements and their variation over time is critical to maintain performance over product life. For example, to maintain operation at the ITU frequency, slight adjustments can be made to a tuning parameter to maintain operation (a constant flat line) at the required optical frequency over life, data not shown [8].

13.6.1 Module level reliability

In addition to the PIC level reliability testing, the Marvell Silicon Photonics modules also go through extended reliability testing at the module level. Fig. 13.13 shows data for the MZM Phase shift for a coherent Si PIC module extended to 5000 hours HTOL. The MZM phase is observed to be stable through the course of the extended test.

For InP PICs with integrated active, passive and sense elements, it is useful to show adjustments applied throughout field operation to quantify changes in component performance through its lifetime. Fig. 13.14A shows the optical frequency from a 100 Gbs (10 Gbs × 10 channel) optical transmitter module change versus field hours for over 15.2 years of continuous operation. The 10 trend lines indicate a slight decrease of a few GHz would have occurred from the laser. The servo

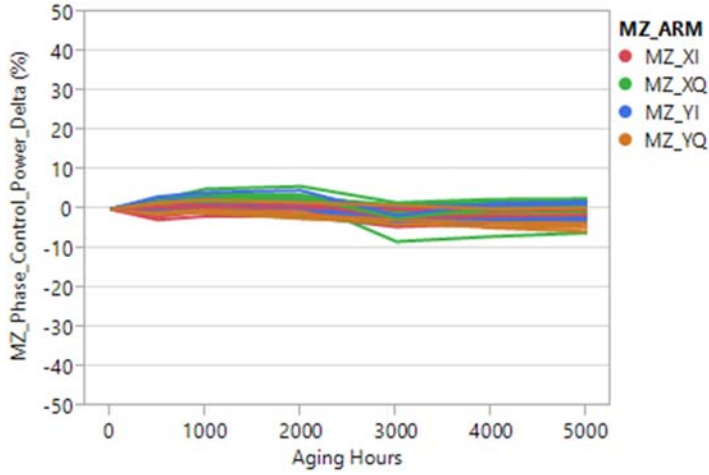

Figure 13.13 MZM phase shift during extended HTOL for silicon photonics coherent module.

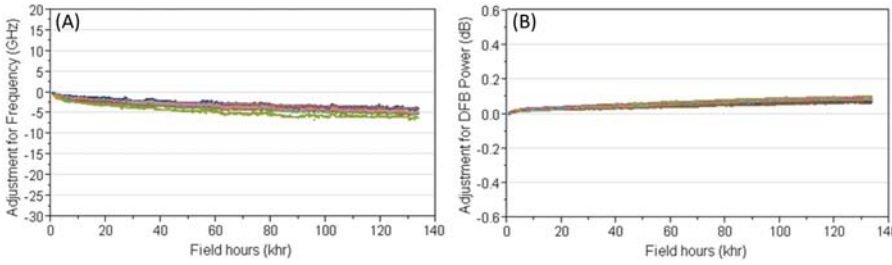

Figure 13.14 (A) Optical frequency from Infinera's 10 channels × 10 Gb/s InP-PIC transmitter module, showing 133,320 h (15.2 years) of operation. Trend indicates a decrease in DFB optical frequency and a trend which slows down over time. All channels are stable to within −6.5 GHz during the lifetime. (B) Optical power throughout the same time for the same module.

adjustments apply exactly the opposite tuning signal to maintain each channel's required optical frequency over life.

Fig. 13.14B shows the corresponding optical power plot for the constant-current-driven DFBs. The monitored power shows a slight increase that slows over the 15.2 years and is within 0.1 dB on all 10 channels of the module. The system uses feedback from the integrated power monitor element to subsequently apply an optical attenuation adjustment to maintain constant optical output power over time. Deployment of stable and predictable optical components is a prerequisite to support a reliable optical network. Fine adjustments from network servo control further guarantee the network operates at optimum performance throughout its life.

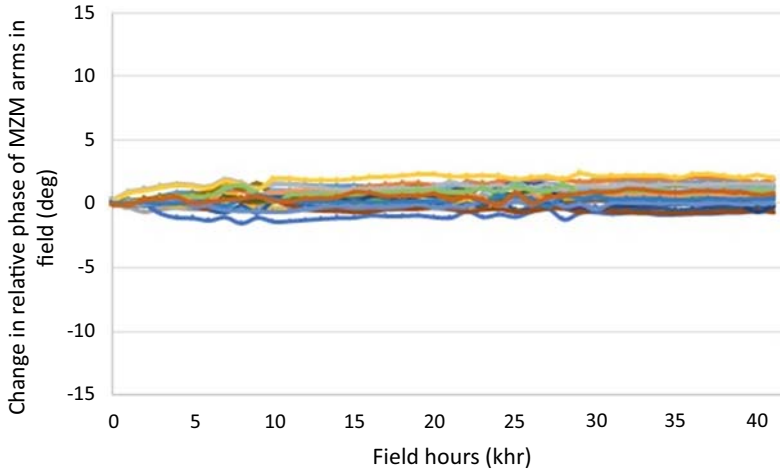

Figure 13.15 Relative phase change from MZM field data, Infinera's 500 Gb/s coherent InP PIC.

Systems with coherent InP PICs (500 Gb/s) have been deployed in a DCI environment since 2013. While coherent modulation offers higher data rates and electronic dispersion compensation, its implementation requires far more components than amplitude modulation as well as phase and amplitude balancing. Each channel has 8 sub-MZM arms and 4 super-MZM arms [3,8]. Here is where photonic integrated circuits truly display benefits. Any arm-pair can be balanced with an electronic signal. Because there are 10 channels (50 Gb/s per channel) integrated on each TxPIC, there would be 20 pairs of super-MZM arms that maintain phase balance. Fig. 13.15 shows field data from the 20 super phase adjuster arm pairs of one of Infinera's 500 Gb/s coherent modules. To maintain a high-quality modulation, the system monitors the constellation diagram and sends a phase adjustment signal to maintain quadrature phase alignment of the relative. This phase adjustment signal is tracked over time during field operation. The 20 traces show that over the 41,000 hours of field operation, only minor phase correction is required.

The relative phase of the InP MZM arms is stable and requires only a few angular degrees of adjustment. Although the phase misalignment is small, by applying minimal corrections the link BER is kept as high as at beginning of life. Prior to field deployment, it has been determined how much margin in either direction would be required to maintain phase alignment in the worst case over the full life of the product.

Building on a generation of 500 Gb/s PIC, Infinera developed the next generation of transceivers using widely tunable lasers (WTL) that can tune to any wavelength within the C-band (and subsequently in the L-band). These lasers serve as the CW optical carrier source in the Tx modules and as the local oscillator in the Rx modules. Each module has 6 wavelengths (a total of 1.2 Tb/s data rate) requiring tracking of 12 optical frequencies per transceiver. Fig. 13.16 shows the field

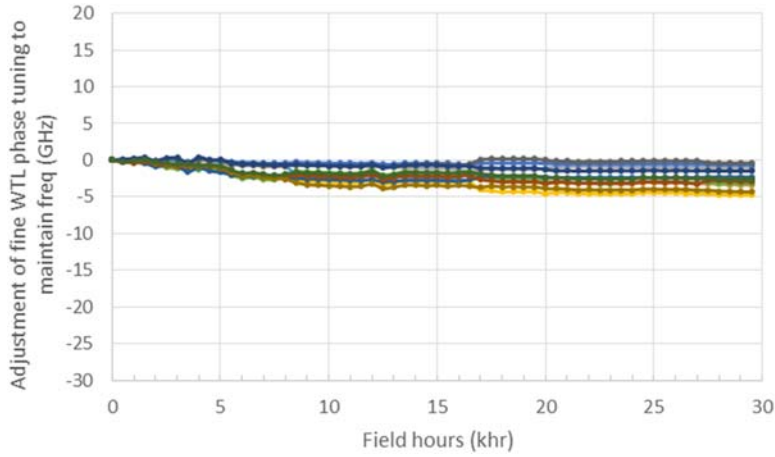

Figure 13.16 Optical frequency servo adjustment for 12 lasers from Infinera's InP 1.2 Tbps Tx/Rx WTLs operating in the field for 30,000 h.

data from 30,000 hours of operation for all 12 WTLs. The optical frequency in the network is maintained within ± 0.3 GHz, these curves represent the projected optical change of the servo-free frequencies of the 12 WTLs. Here, a slight and saturating decrease in effective index (in this case for the fine-tuning WTL phase section) was typically applied by the system to maintain precise optical frequencies for the WTLs. Because sufficient tuning margin has been set aside at beginning of life, these minor changes enable maintaining a constant optical frequency over the lifetime (Figs. 13.15 and 13.16).

The InP module field data presented summarize the adjustments to individual components. However, the net failure rate takes into account all the individual components of the PIC. Infinera's 100-Gb/s PICs pairs have accumulated over 3 billion field hours with zero field failures, giving a FIT rate of <0.3 at 60% confidence level [8,9].

The PIC and module field performance examples demonstrated that by following a DFQ practice and building robust models it is possible to achieve a reliable optical transceiver product that exceeds the lifetime requirements of DC and HPC modules.

13.7 Strategy for robust manufacturing quality

The combined need for consistent reliability and high-volume manufacturing greatly raises the requirement for detecting and screening parts with any meaningful deviation in the fabrication process. Wafers need to be evaluated beyond the normal performance screens to ensure acceptable pedigree. Performance screens are effective for flagging yield issues, but tracking the raw material and fab process history

can provide more rapid feedback and is able to flag deviations that may still be within product specification yet can indicate a process deviation that impacts over-life stability.

13.7.1 Wafer qualifications/certification

In addition to the standard product qualification at the initial production release stage, there may be certain scenarios where a sampling reliability qualification is done on every single wafer or wafer fabrication lot. For elements where 100% burn-in is not done or required, the per wafer or per wafer lot reliability sampling is used to ensure the reliability of the product on an on-going basis. This is designated as wafer qualification. In other cases, designated as wafer certification, the wafer or lot sampling can be used to decide whether the yield economics justify proceeding with the 100% burn-in of the entire wafer or lot. Typically, this wafer or lot sampling is done as a biased aging test (accelerated aging, HTOL etc.) but can also be done as an unbiased aging test (High Temperature Storage, Temperature Cycling, Damp Heat testing etc.)

To supplement the reliability feedback and assure thorough evaluation for the more stringent, rapid demands of DCI and HPC applications, it is beneficial to periodically run sampled wafer evaluation cells to confirm that stability of the process is maintained. Although a wafer fab process is frozen and qualified, one needs to account for variations over time due to incoming material, supplier specs and fabrication equipment to maintain a reliable supply of production PICs. This can be achieved by on-going periodic evaluation of key parameters beyond the standard process screens. These evaluations are performed under accelerated conditions to verify epitaxial growth quality, the stability of metallization and the development of mechanical stress. The tests might be executed directly on wafer or in quick-turnaround, die-based configurations. These type of wafer certifications give high confidence that the fabrication process continues to be in-line with the initial qualification wafers.

13.7.2 High volume manufacturing

For high volume manufacturing, to ensure reliable component and supply chain continuity, the manufacturing process must be well monitored and controlled. Silicon photonics is no exception and has the advantage of benefitting from the extensive set of metrology tools available to the silicon wafer processing world to monitor film thicknesses, etch rates, critical dimensions (CD) and other material properties like dielectric constant and refractive index. Silicon photonics uses advanced 8″ and 12″ silicon toolsets. However, silicon photonics differs from standard CMOS because parameters important for optical performance are not necessarily the same as those for electrical performance of a CMOS process. An example, small changes in the silicon waveguide width can result in several tenths of a dB of optical loss. Hence, in addition to controlling the photolithography, the etch rate and etch bias of silicon must be controlled as well. The resulting width of the

silicon waveguide is measured and monitored as a critical dimension (CD) with standard statistical process control methodologies to ensure the optical loss is well controlled. Similarly, silicon nitride thickness, refractive index and CD are monitored to ensure the optical loss of the nitride waveguides is well controlled.

In addition to in-process monitoring, both electrical and optical measurements are made on every silicon photonics wafer to ensure that the wafers meet specifications and will have high-yielding product die. Besides the normal electrical checks for metal and implanted Si sheet resistivities, metal and via open/shorts monitors and junction capacitances typically measured for CMOS wafers, unique to the silicon-photonics wafers, we measure photodiode (PD) dark current, responsivity, PD capacitance, Mach-Zehnder Modulator (MZM) turn-on voltage, MZM capacitance and more at wafer completion. To further monitor quality, arrays of devices are often used to obtain a more statistical outlook of the device health. Wafers are dispositioned against a wafer acceptance criteria (WAC) which a given number of tested sites must pass for the wafer to be shipped to the customer.

In summary, beyond the normal process controls for CMOS processes, it is important to monitor additional parameters on silicon photonics wafers to ensure good optical performance and manufacturing control. Most of these can be measured with standard fab metrology tools but do require an optical test system to measure optical properties of passives and active devices for wafer disposition. In doing so, we can ensure that only high quality, reliable wafers will be delivered to the end customer.

The importance of monitoring process parameters at both wafer fab and assembly processes cannot be overemphasized. The fabrication of InP PICs requires monolithic integration of a wide range of functions and components, such as, lasers, optical amplifiers, modulators and detectors, and precise wavelength control over multiple InP alloyed layers (bandgaps) with different doping types and levels. Process monitors at both wafer fab and assembly processes serve as performance and reliability indications as changes that lead to performance and reliability impacts are intimately linked. Consequently, statistical process controls (SPC) are used routinely in an InP PIC commercial line to analyze and control complex epitaxy and fabrication process. Infinera's PIC process is set up to automatically track well over 100 parameters as they can have meaningful impact on transceiver reliability. The list includes epitaxial, metal and dielectric thickness and stresses, photolithography critical dimensions, optical index and absorption, isolation resistances, device performance (such as PD leakage current and responsivity, laser frequency, slope efficiency/tuning rate), visual defects and burn-in (BI) shifts. Any tracked parameter that exhibits changes beyond statistically expected limits may indicate a reliability deviation and requires an FMEA, and perhaps derisking, and, if significant would require a set of qualification cells.

13.7.3 On-going reliability monitoring (PIC and module)

Beyond periodic wafer-level sampling and monitoring, there may also be changes in reliability due to steps subsequent to wafer fabrication including: die singulation,

antireflection coating, die attach, wirebonding, chip-on-carrier testing, burn-in, handling, module assembly, and module test. Cleaving and coating are known to be critically important to reliability on many types of edge emitting lasers, such as the AlGaInAs lasers that are popular for 1310 nm uncooled transceivers. Therefore it is beneficial to periodically run sampled aging cells on an on-going basis at both the PIC and module level to assure reliability is equivalent to the initial product qualification.

Typically, the ongoing reliability monitoring test (ORM or ORT), is performed monthly or quarterly, depending on the volume of the product. These tests should sample the production supply from the current month or quarter to evaluate stability of key parameters in accelerating aging cells. The ORM plan includes the key biased and unbiased tests that are done as part of the initial qualification and should sample multiple wafers and lots. The sample sizes for these tests are usually smaller than the initial qualification. These periodic ongoing accelerated aging cells are in addition to the initial qualification and in addition to the standard process screens which would apply to every part.

In addition to the reliability testing that is done to ensure ongoing product quality, a closed-loop FA process for continuous improvement is critical. This includes FA for rootcause and corrective action implementation all production line and ORM failure modes as well as customer returns. At Marvell, the closed-loop FA process has been instrumental in reducing the field return rate for Si Photonics PIC modules in high volume manufacturing (HVM) by over 5× compared to the early new production introduction phase of the product as seen on Fig. 13.17. This plot presents the normalized cumulative dppm (defect rate, failure rate) versus shipment volume demonstrating a reduction in failure rates with volume compared to early designs/process parts.

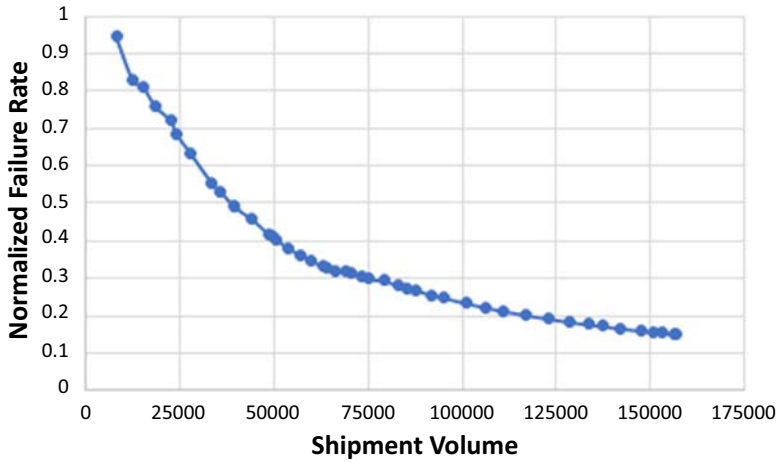

Figure 13.17 Normalized field return rate for Marvell SiPh PIC modules reduced by over 5× with closed-loop failure analysis for product improvement.

13.8 Challenges and opportunities

We have presented data to support the reliable field performance of photonic ICs, based on SiP and InP platforms using established design for quality methodologies. The deployment of photonic ICs into DCI and HPC applications faces a number of challenges in reliability in the near future. One challenge is that of the density of connections which will be required. As with Moore's Law, which requires transistor reliability to improve on a per-transistor basis as transistor integration increases, so too the number of lasers and parallel links per box has increased in many of the new generations. Another challenge is increasing optical power required to support higher baud rates with directly modulated lasers, which means operating the lasers with higher bias current, thereby reducing the lifetime of the lasers. Some of these PIC challenges have been addressed in telecommunication and the learnings can be ported to DCI and HPC applications. However, the GR-468 standards, which call for only 22 devices to be tested, are no longer providing adequate assurance of reliability. Documented cases have been shown where designs passed the GR-468 issue 2 spec with no failures but gave unacceptable field failure rates of up to 2000 FIT later in large cloud data center deployments [39,40]. This has driven a desire to update the standard to increase the requirements, both by increasing the quantities required, as well as by adding more specific lot or wafer qualification requirements.

13.8.1 Opportunities for new validation tools/technology

Improved methods for reliability testing have helped to reduce the cost of data collection on PICs. For example, wafer-level burn-in has been available for nearly 20 years on simple VCSEL wafers, but wafer-level PIC burn-in and aging has only become popular more recently, and enables testing to increase by at least an order of magnitude for the same cost profile, since the costly packaging of thousands of devices can be skipped if the inherent component reliability is the object that we wish to study.

13.8.2 Challenges with emergent technologies

"Embedded Optics" is another upcoming challenge for DC/HPC applications. Today's fiber optic transceivers are "faceplate pluggable," meaning that if maintenance is required, they can be easily exchanged for replacement transceivers in just a few minutes. But "embedded optics" will be deployed inside the switching box and require pulling the entire switch out of operation to replace any failed links. Embedded optics, and increased density, mean higher deployment temperatures than were observed in the previous faceplate pluggable optics. Temperatures as high as 110°C are cited in embedded optics deployments. In some cases, the optoelectronic links are socketed, while in others they are soldered in place, experiencing a different environment than has previously been reported. As was the case with the telecom industry, new designs being considered include redundant

channels which can act as backups, and can either dynamically switch to the backup laser, or be switched with a "software repair."

13.8.3 Opportunities across PIC capability and assembly improvements

Assembly methods continue to evolve, including more chip stacking, bump bonding of PICs directly onto ICs, and nonhermetic designs. Bending stresses on the PICs can be significant, and cause reliability problems, unless care is taken to control it in the design and assembly process.

As the various applications move to higher density, high baud rate and higher power dissipation density, we expect those will also drive evolution of integrated optics systems. An example of this is for instance the push for closer integration of optical transceiver functionality with ethernet switch ASICs, as discussed in Chapter 8. This is driven by the attempt to reduce overall power dissipation by reducing the electrical power dissipated in the interconnect between ASICs and transceivers. Similar is the demand for closer integration of optical I/O to memory and compute nodes in computing applications, the reader is referred to Chapter 9 for a detailed discussion. Those "copackaged" applications of optics and electronics will drive even more technology elements, such as high-speed photonics (100 G/lane, 200 G/lane), with low parasitic interconnect between optical and electrical. Those systems are also expected to operate integrated optics assemblies and packaging elements (e.g., fiber arrays, optical bonds and optical connectors) at higher temperatures ($\sim$85°C−105°C).

We believe that this evolution will push for the development and introduction of advanced packaging technologies, as referenced in Chapter 12, and development of improved material systems for higher performance and higher stability at more demanding environmental conditions.

References

[1] F.R. Nash, et al., Implementation of the proposed reliability assurance strategy for an InGaAsp/InP, planar mesa, buried heterostructure laser operating at 1.3 μm for use in a submarine cable, ATT Tech J. 64 (3) (1985) 809−860. Available from: https://doi.org/10.1002/j.1538-7305.1985.tb00449.x.

[2] GR-468-CORE, Issue 2, September, 2004: Generic Reliability Assurance Requirements for Optoelectronic Devices Used in Telecommunications Equipment, https://telecom-info.njdepot.ericsson.net/site-cgi/ido/docs.cgi?ID = SEARCH&DOCUMENT = GR-468 (accessed March 15, 2023).

[3] F. Kish, et al., From visible light-emitting diodes to large-scale III−V photonic integrated circuits, Proc. IEEE 101 (10) (2013) 2255−2270. Available from: https://doi.org/10.1109/JPROC.2013.2275018.

[4] Y.A. Akulova, et al., Widely tunable electroabsorption-modulated sampled-grating DBR laser transmitter, IEEE J. Sel. Top. Quantum Electron. 8 (6) (2002) 1349−1357. Available from: https://doi.org/10.1109/JSTQE.2002.806677.
[5] J.S. Barton, E.J. Skogen, M.L. Masanovic, S.P. Denbaars, L.A. Coldren, A widely tunable high-speed transmitter using an integrated SGDBR laser-semiconductor optical amplifier and mach-zehnder modulator, IEEE J. Sel. Top. Quantum Electron. 9 (5) (2003) 1113−1117. Available from: https://doi.org/10.1109/JSTQE.2003.819477.
[6] R. Nagarajan, et al., Large-scale photonic integrated circuits, IEEE J. Sel. Top. Quantum Electron. 11 (1) (2005) 50−65. Available from: https://doi.org/10.1109/JSTQE.2004.841721.
[7] R.A. Hyle, Fiber optics-failure modes and mechanisms, in: Annual Reliability and Maintainability Symposium 1992 Proceedings, 1992. Available from: https://doi.org/10.1109/ARMS.1992.187852.
[8] F. Kish, et al., System-on-chip photonic integrated circuits, IEEE J. Sel. Top. Quantum Electron. 24 (1) (2018) 1−20. Available from: https://doi.org/10.1109/JSTQE.2017.2717863.
[9] G.E. Hoefler, et al., Foundry development of system-on-chip InP-based photonic integrated circuits, IEEE J. Sel. Top. Quantum Electron. 25 (5) (2019) 1−17. Available from: https://doi.org/10.1109/JSTQE.2019.2906270.
[10] L. Li, J. Xue, Reliability of microelectronics packaging in the era of EnergyWise and Borderless networks, in: Reliability Physics Symposium (IRPS), 2010 IEEE International, 2010, pp. 440−445. Available from: https://doi.org/10.1109/IRPS.2010.5488791.
[11] While not driven by standards requirements, a 500 parts-per-million defect level ('500 dppm' or 1/20th of a percent) is a widely held expectation for quality level for fiber optic transceivers, and is the goal for a wide range of suppliers and customers. It has proven to be a good balance between cost of quality on the producer side, and cost of maintenance on the customer side.
[12] D.H. Stamatis, Failure Mode and Effect Analysis: FMEA from Theory to Execution, second ed., ASQC Quality Press, Milwaukee, WI, 2003.
[13] MIL-STD-1629 A Procedures Performing A Failure Mode. http://everyspec.com/MIL-STD/MIL-STD-1600-1699/MIL_STD_1629A_1556/ (accessed July 16, 2021).
[14] Potential Failure Mode and Effects Analysis (FMEA) | JEDEC. https://www.jedec.org/standards-documents/docs/jep131b (accessed July 16, 2021).
[15] P.A. Tobias, D. Trindade, Applied Reliability, third ed., Chapman and Hall/CRC, 2011.
[16] S. Ross, First Course in Probability, ninth ed., Pearson, 2012.
[17] W.Q. Meeker, L.A. Escobar, Pitfalls of accelerated testing, IEEE Trans. Reliab. 47 (2) (1998) 114−118. Available from: https://doi.org/10.1109/24.722271.
[18] A. Narasimha, et al., A 40-Gb/s QSFP Optoelectronic Transceiver in a 0.13 μm CMOS Silicon-on-Insulator Technology, in: OFC/NFOEC 2008—2008 Conference on Optical Fiber Communication/National Fiber Optic Engineers Conference, February 2008, pp. 1−3. Available from: https://doi.org/10.1109/OFC.2008.4528356.
[19] T. Pinguet, B. Analui, E. Balmater, D. Guckenberger, G. Capellini, Monolithically Integrated High-Speed CMOS Photonic Transceivers, 2008. Available from: https://doi.org/10.1109/GROUP4.2008.4638200.
[20] X. Ding, D. Wang, L. Zhang, L. Chen, W. Zhang, Kirkendall voids evolution in Sn-Ag-In solder bumps and its effect on reliability, in: 2016 17th International Conference on Electronic Packaging Technology, 2016, Available from: https://doi.org/10.1109/ICEPT.2016.7583328.
[21] M.D. Cheng, S.Y. Chang, S.F. Yen, T.H. Chuang, Intermetallic compounds formed during the reflow and aging of Sn-3.8Ag-0.7Cu and Sn-20In-2Ag-0.5Cu solder ball grid

array packages, J. Electron. Mater. 33 (3) (2004) 171−180. Available from: https://doi.org/10.1007/s11664-004-0176-9.
[22] R. Labie, et al., Outperformance of Cu pillar flip chip bumps in electromigration testing, in: 2011 IEEE 61st Electronic Components and Technology Conference ECTC, 2011, Available from: https://doi.org/10.1109/ECTC.2011.5898531.
[23] H. Kitada, H. Tashiro, S. Miyahara, A. Dote, S. Tadaki, S. Sakuyama, Thermal stress reliability of copper through silicon via interconnects for 3D logic devices, in: 2016 IEEE 18th Electronics Packaging Technology Conference (EPTC), Singapore, Singapore, November 2016, pp. 115−119. Available from: https://doi.org/10.1109/EPTC.2016.7861455.
[24] C. Cassidy, et al., Through silicon via reliability, IEEE Trans. Device Mater. Reliab. 12 (2) (2012) 285−295. Available from: https://doi.org/10.1109/TDMR.2012.2189212.
[25] G. Hariharan, et al., A comprehensive reliability study on a CoWoS 3D IC package, in: 2015 IEEE 65th Electronic Components and Technology Conference (ECTC), San Diego, CA, May 2015, pp. 573−577. Available from: https://doi.org/10.1109/ECTC.2015.7159648.
[26] H.-N. Huang, Y.-C. Liao, W.-T. Tseng, C.-T. Lin, C.-H. Chiu, Study of underfill material for fine pitch Cu pillar bump, in: 2011 6th International Microsystems, Packaging, Assembly and Circuits Technology Conference IMPACT, 2011, Available from: https://doi.org/10.1109/IMPACT.2011.6117231.
[27] M. Fukuda, Reliability and Degradation of Semiconductor Lasers and LEDs, Artech House, Boston, 1991.
[28] P.W. Epperlein, Semiconductor Laser Engineering, Reliability and Diagnostics: A Practical Approach to High Power and Single Mode Devices, first ed., Wiley, 2013. Available from: http://doi.org/10.1002/9781118481882.
[29] MIL-STD-883 Test Methods. pp. 1004−1018. https://eesemi.com/milmethods.htm (accessed September 11, 2021).
[30] H. Greenhouse, Hermeticity of Electronic Packages, Noyes Publications, Park Ridge, NJ/Norwich, NY, 2000.
[31] A.W. Fang, H. Park, O. Cohen, R. Jones, M.J. Paniccia, J.E. Bowers, Electrically pumped hybrid AlGaInAs-silicon evanescent laser, Opt. Express 14 (20) (2006) 9203−9210. Available from: https://doi.org/10.1364/OE.14.009203.
[32] R. Jones, P. Doussiere, J. Driscoll, W. Lin, Y. Akulova, T. Komljenovic, J. Bowers, Heterogeneously integrated InP / silicon photonics, IEEE Nanotechnology Magazine 13 (2) (2019) 17−26. Available from: https://doi.org/10.1109/MNANO.2019.2891369. In press.
[33] R. Blum, Integrated silicon photonics for high-volume data center applications, in: Optical Interconnects XX, San Francisco, United States, February 2020, p. 19. Available from: https://doi.org/10.1117/12.2550326.
[34] Z. Zhang, J. Liu, P. Zhao, G.Z. Xiao, C.P. Grover, Active alignment of optical fibers to planar waveguides using a thermal-curing adhesive, J. Light. Technol. 23 (2) (2005) 567.
[35] S. Mitachi, A. Hagiwara, N. Murata, H. Kageyama, K. Kimura, Bond properties of silane modified optical adhesives for silicon photonics, in: 2014 International Conference on Electronics Packaging (ICEP), April 2014, pp. 712−715. Available from: https://doi.org/10.1109/ICEP.2014.6826772.
[36] M.A. Uddin, et al., Delamination problems of UV-cured adhesive bonded optical fiber in V-groove for photonic packaging, IEEE Photon. Technol. Lett. 16 (4) (2004) 1113−1115. Available from: https://doi.org/10.1109/LPT.2004.824639.

[37] M.A. Uddin, H.P. Chan, T.O. Tsun, Y.C. Chan, Uneven curing induced interfacial delamination of UV adhesive-bonded fiber array in V-groove for photonic packaging, J. Light. Technol. 24 (3) (2006) 1342−1349. Available from: https://doi.org/10.1109/JLT.2005.863329.

[38] V. Lal, et al., Extended C-band tunable multi-channel InP-based coherent transmitter PICs, J. Light. Technol. 35 (7) (2017) 1320−1327. Available from: https://doi.org/10.1109/JLT.2017.2670547.

[39] A. Chakravarty, S. Giridharan, M. Kelly, A. Poojary, V. Zeng, Characterizing large-scale production reliability for 100G optical interconnect in facebook data centers, Proceedings, Frontiers in Optics, 2017 1 (2017) JW4A.64. Available from: https://doi.org/10.1364/FIO.2017.JW4A.64.

[40] V. Zeng, Practical Reliability Guidance for Laser in PIC Manufacturing, 2020. Available from: https://lnkd.in/gnfidYv (accessed 15.03.23).

Index

Note: Page numbers followed by "*f*" and "*t*" refer to figures and tables, respectively.

A
Active devices, 343–344
Active optical cables (AOCs), 4, 233
Active–passive $8 \times 8 \times 8\lambda$ semiconductor optical amplifier space, 282–285, 284*f*
Active–passive 16×16 semiconductor optical amplifier space switch, 281–282, 282*f*
Adiabatic couplers, 179
Advanced Ge photodiode processing on CMOS platform, 136–142, 137*f*
  advanced integration of Ge devices with CMOS electronics, 139–142
  processing challenges with Ge photodetectors, 137–139
All-active 16×16 semiconductor optical amplifier space switch, 278
All-to-all topologies, 6–7
Aluminum copper (AlCu), 379
AMDs xGMI, 3
Amplifiers, 344
Antireflection (AR), 402
AOC connections, 12
AOCs. *See* Active optical cables (AOCs)
AOP. *See* Average optical power (AOP)
Application-specific integrated circuits (ASICs), 267–268, 345–346
AR. *See* Antireflection (AR)
Arbitrary waveform generator (AWG), 403
Arrayed waveguide gratings (AWG), 176, 183–185, 184*t*, 282–283, 372
Arrayed waveguide-grating router (AWGR), 326–327
Arsenic (As), 137
Artificial intelligence workloads, 3–4, 18–19
ASICs. *See* Application-specific integrated circuits (ASICs)
Autonomous vehicles, 131
Avalanche photodetectors, 55–56, 142–149
  avalanche photodiodes (APDs)
    and Ge/Si avalanche photodiode, 128, 143–147
    next-generation applications of, 147–149
  breakthrough and current market, 142–143
Average optical power (AOP), 239
AWG. *See* Arbitrary waveform generator (AWG); Arrayed waveguide gratings (AWG)
AWGR. *See* Arrayed waveguide-grating router (AWGR)

B
Back end of line (BEOL), 107, 127, 149, 379
Back-annotation, 351
Bandwidth (BW), 87–90, 233–234, 290
Banyan-type blocking networks, 270–271
Banyan-type switch fabric, 270–271
Barium titanate ($BaTiO_3$), 76, 107, 243–244
  on Si, 105–108
BCB layers. *See* Benzocyclobutene layers (BCB layers)
Beam propagation method (BPM), 341–342, 372
Beam steering, 77–78
Beam width, 417
Bends, 164, 372–373
Beneš networks, 270
64×64 Benes switch, 278–279
Beneš topology, 271–273
Benzocyclobutene layers (BCB layers), 41–42

BEOL. *See* Back end of line (BEOL)
BEP. *See* Bidirectional eigenmode propagation (BEP)
BER. *See* Bit error ratio (BER)
BI. *See* Burn-in (BI)
Bidirectional eigenmode propagation (BEP), 341
Bidirectional encoder representations from transformers (BERT), 3–4
BIST. *See* Built-In Self-Test (BIST)
Bit error ratio (BER), 11–12, 314–317
Bitrate and 3-dB bandwidth change of MZMs, 96*f*
Bonding process, 48
Boron (B), 137
Boutique technology, 320
BOX. *See* Buried oxide (BOX)
BPM. *See* Beam propagation method (BPM)
Bragg GCs. *See* Bragg grating couplers (Bragg GCs)
Bragg grating couplers (Bragg GCs), 244–245
Breakdown voltage, 144
Built-In Self-Test (BIST), 427
Buried oxide (BOX), 139, 379
Burn-in (BI), 452
Butt-coupling, 242, 376–377
Butterfly, 9–10
BW. *See* Bandwidth (BW)

C

CAGR. *See* Compounded annual growth rate (CAGR)
Capacitance-loaded traveling-wave electrode (CL-TWE), 94–96
Carrier accumulation, 73
Carrier depletion, 73
Carrier injection, 73
Carrier transport process, 128
CD. *See* Critical dimensions (CD)
CDC. *See* Contra-DC with Bragg grating (CDC)
CDR. *See* Clock and data recovery (CDR)
CE. *See* Coupling efficiency (CE)
Chemical mechanical polishing (CMP), 359, 363
Chemical vapor deposition (CVD), 132
  epitaxial growth of Ge on Si by, 132
Chemo-mechanical polishing (CMP), 132
Chip-level end-of-line optical characterizations, 376–377
Chip-on-chip-on-substrate (CoCoS), 451–452
Chromatography, 447
Circuit(s), 159
  design, 337
  simulation, 349–350
    interface, 350
  topology, 90–92
CL. *See* Confidence level (CL)
CL-TWE. *See* Capacitance-loaded traveling-wave electrode (CL-TWE)
Clock and data recovery (CDR), 235–236
Clos networks, 5, 270
Cloud computing, 2–3
Cloud-scale providers (CSP), 201–202, 316
CMOS process, 44
  advanced Ge photodiode processing on CMOS platform, 136–142
  advanced integration of Ge devices with CMOS electronics, 139–142
    photonic device integration with silicon CMOS process, 141*f*
  circuitry, 392
  compatible processing, 53
CMP. *See* Chemical mechanical polishing (CMP); Chemo-mechanical polishing (CMP)
Coarse wavelength division multiplexing (CWDM), 42, 175, 246–248, 394
COBO. *See* Consortium for On-Board Optics (COBO)
CoCoS. *See* Chip-on-chip-on-substrate (CoCoS)
Coefficient of thermal expansion (CTE), 134–135, 440, 457–458
Coherent communication schemes, 336
Coherent interconnects for data centers
  400ZR deployment, 227*f*
  coherent dense wavelength-division multiplexing transceivers, 227*f*
  DCI, 201–203
  development paths for volume deployment, 219–224
    market impact of IPoDWDM shift, 223–224
    OIF 400ZR specifications, 221

role of standards, consortia, and MSAs, 219–220
ZR + derivatives, 221–223
direct detection and coherent links for DCI, 203–210
interdatacenter interconnects, 225–226
link efficiency projections as data rates, 228*f*
next in standards, 226–227
photonic-integrated circuits for coherent DCI, 210–219
Coherent links for DCI, direct detection and, 203–210
Coherent optical devices, 205
Coherent PICs, potential for further cost, size, and power reductions for, 215–219
Coherent receivers, 213
Coherent technologies, 16
Coherent transceivers, 210–214
first I–Q modulator, 211*f*
InP coherent modulator with integrated optical gain, 212*f*
plots of transmission through Mach–Zehnder modulator *vs.* differential phase, 211*f*
progression of integrated coherent optics, 213*f*
thin-film $LiNbO_3$ coherent modulator, 212*f*
Coherent transmission formats, 260
"ColorZ" modules, 206–207
Comb lasers, 39–41
Kerr comb lasers, 40–41
mode-locked laser, 40
Combiners, 168–170
Commodity interconnects, 7–8
Compounded annual growth rate (CAGR), 322
"Compression" algorithms, 131
Compute infrastructure, 1
Computer-aided design, 371–372
Confidence level (CL), 440
Connector coupler, 416
Consortium for On-Board Optics (COBO), 21
Constrained topologies, 6–7
Continuous wave DFB (CW DFB), 249
Continuous wavelength tuning, 39
Continuous-wave (CW), 203–204
Contra-DC with Bragg grating (CDC), 176–177
Contra-directional Bragg grating filters, 183
Control planes, scheduling, 322–324
Control power consumption, 295
Conventional coherent architectures, 225
Conventional III–V modulators, 93*f*, 94–96, 96*f*
Conventional planar fabrication techniques, 104
Co-packaged optics (CPO), 16, 29, 42, 214, 236–237, 309, 315
integration to support scaling to 400-to 800-G transceivers, 215*f*
modules, 21
scaling from 25T-100T network switches, 215*f*
and switch assembly, 27*f*
Co-packaging, 25, 417–418
for data center interconnects, 254–257
of photonic device and electronic ICs, 419*f*
Copper (Cu), 378–379
Corner-based analysis, 361
Correlated process variations, 361
"Cosmic Cube", 5
Couplers, 107, 168–170
Coupling efficiency (CE), 395, 414–415
CPO. *See* Co-packaged optics (CPO)
Cray SlingShot, 4
Cray XT series systems, 5
Critical dimensions (CD), 422–423, 463–464
Cross-linking process, 456
Cross-section of 3D-integrated silicon photonics CoCoS assembly, 451*f*
Crossbar networks, 270, 273
CSP. *See* Cloud-scale providers (CSP)
CTE. *See* Coefficient of thermal expansion (CTE)
Current modes, 90–91
CVD. *See* Chemical vapor deposition (CVD)
CW. *See* Continuous-wave (CW)
CW DFB. *See* Continuous wave DFB (CW DFB)
CWDM. *See* Coarse wavelength division multiplexing (CWDM)

D
2.5D integration, 414, 419–420
DACs. *See* Direct-attach copper cables (DACs)
Damascene process, 132
Dark current (Idk), 124, 446
Data center interconnect (DCI), 63, 201–203, 233–241
  and communications, 35
  DCI-targeted coherent chipset, 220
  direct detection and coherent links for DCI, 203–210, 203*f*, 204*f*, 205*f*
    coherent optical devices, 205
    evolution of DCI link technologies from intensity modulated direct detection to coherent, 206–210
  electrical signal integrity, 237–239
  evolution of DCI link technologies, 206–210, 207*f*, 208*f*, 209*f*
  integrated photonic building blocks for, 241–245
  intradata center link requirements and standards, 234–237
  photonic integration/copackaging for, 254–257
  photonic-integrated circuit loss bottlenecks and challenges, 239–241
  regional datacenter architectures, 202*f*
  reliability of DCI optical systems and field data, 457*f*, 458–462
  segmentation of datacenter links by distance, 201*f*
Data center network (DCN), 265
Data center reach (DR), 233–234
Data centers (DCs), 1, 149, 171, 233, 266–267, 288–294, 439
  bandwidth and transceiver growth, 234*f*
  DC/HPC networks, 265–266
  optical interconnect requirements, 449–450
  photodetectors in, 123–132, 149
  reliability requirements for DC/HPC application environment, 441–443, 442*t*
  reliability standards for DC and HPC applications, 443–444
Data communications, 53
  interconnect figures of merit, 23–24
  major inflection points and challenges, 24–30
  optical module form factors, 19–23
  optical network case studies, 1–19
    enterprise data center, 1–8
    high-performance computing, 1–8
    hyperscale data centers, 2–3
Data parallelism, 316
Data planes, scheduling, 322–324
Data-parallel machine learning training (Data-parallel ML training), 3
Data-parallel ML training. *See* Data-parallel machine learning training (Data-parallel ML training)
Datacom
  interconnects, 19
  types of sources needed for, 35–41
    comb lasers, 39–41
    multiwavelength lasers, 36–38
    single frequency lasers, 38
    tunable lasers, 38–39
3-dB bandwidth, 124
– 3 dB electro-optical bandwidth (BW_{3dB}), 70–71
DBR lasers. *See* Distributed Bragg reflector lasers (DBR lasers)
DCI. *See* Data center interconnect (DCI)
DCN. *See* Data center network (DCN)
DCs. *See* Data centers (DCs)
DD. *See* Direct detection (DD)
DDL. *See* Distributed deep learning (DDL)
DDM. *See* Digital diagnostics monitoring (DDM)
Deep learning (DL), 288
Demultiplexer (Demux), 140, 175–185
Demux. *See* Demultiplexer (Demux)
Dense wavelength division multiplexing (DWDM), 40, 65, 201
Depletion modulators, 78–80
Derisking plan, 445–448, 448*f*
Design and testing, future directions in, 186
Design flow, 347–348, 348*f*
Design FMEA (DFMEA), 445
Design for manufacturability (DFM), 337, 352, 358–364, 440, 444
  manufacturing variation-aware photonic models, 359–361
  manufacturing variations, 358–359
  photonics design
    rule checking, FILL, and optical proximity correction, 363–364

for test and design for reliability, 364
yield analysis and optimization, 361–363
Design for reliability, 364
Design for test (DFT), 347, 352
Design process, 335, 338
Design rule checks (DRC), 347, 371–373
Design-for Quality (DFQ), 440, 444
Detectors, 107, 344
Device design, 337–338
Device metrics, 268–270
Device under test (DUT), 293–294, 424–425
DF. *See* Dragonfly (DF)
DFB lasers. *See* Distributed-feedback lasers (DFB lasers)
DFM. *See* Design for manufacturability (DFM)
DFMEA. *See* Design FMEA (DFMEA)
DFQ. *See* Design-for Quality (DFQ)
DFT. *See* Design for test (DFT)
DGD. *See* Differential group delay (DGD)
Die-to-wafer approach, 388–389
Differential group delay (DGD), 279
Differential Scanning Calorimetry (DSC), 456
Differential-quadrature-phase-shift keying (DQPSK), 211
Digital diagnostics monitoring (DDM), 442–443
Digital signal processor (DSP), 43, 203–204
Direct detection (DD), 101
Direct detection and coherent links for DCI, 203–210
Direct modulated laser (DML), 453–454
Direct-attach copper cables (DACs), 12
Directional couplers, 179
Directly modulated lasers (DMLs), 93–94, 233–234, 248–249, 389–390
Directly modulated multimode VCSELs, 93–94
Discrete lasers, 51–52
Distributed Bragg reflector lasers (DBR lasers), 38, 400
Distributed deep learning (DDL), 288
training based on optical switching, 288–290
Distributed-feedback lasers (DFB lasers), 38, 93–94, 249, 400
DL. *See* Deep learning (DL)
DLG. *See* Double-layer graphene (DLG)
DML. *See* Direct modulated laser (DML)
DMLs. *See* Directly modulated lasers (DMLs)
Dopant diffusion, 137–138
Double-layer graphene (DLG), 102
DPSK, 336
DQPSK. *See* Differential-quadrature-phase-shift keying (DQPSK)
DR. *See* Data center reach (DR)
Dragonfly (DF), 5–10
DRC. *See* Design rule checks (DRC)
Driver, codesign of, 92–93
Drude–Lorentz model, 71–73
Dry etching, 402
Dry lithography patterning process, 393–394
DSC. *See* Differential Scanning Calorimetry (DSC)
DSP. *See* Digital signal processor (DSP)
DUT. *See* Device under test (DUT)
DWDM. *See* Dense wavelength division multiplexing (DWDM)

E

E-Beam lithography process, 167
E/O codesign. *See* Electronic–photonic codesign (E/O codesign)
EA. *See* Electro-absorption (EA)
EAM FKE-based devices, 86
EAMs. *See* Electro-absorption modulators (EAMs)
Echelle diffraction gratings (EDGs), 183
Mux/Demux, 183–185, 184*t*
Echelle grating (EG), 176
ECN. *See* Explicit Congestion Notification (ECN)
EDA. *See* Electronic design automation (EDA)
EDC. *See* Enterprise data center (EDC)
Edge coupling, 52, 172–173, 415–416
EDGs. *See* Echelle diffraction gratings (EDGs)
EG. *See* Echelle grating (EG)
EICs. *See* Electrical integrated circuits (EICs)
Eigen mode expansion (EME), 341
and propagation, 341

Electrical Baseline, 312–313
Electrical cables, 12
Electrical integrated circuits (EICs), 236–237, 343–344, 359, 419*f*, 422, 450
Electrical integration, levels of, 214
Electrical interfaces, 24
Electrical interposer, 417
Electrical packaging, 417–420, 418*f*, 419*f*
Electrical power dissipation, 285
Electrical signal integrity, 237–239, 237*f*, 238*t*
Electrical simulation, notes on, 353–355
Electrical wire bonds, 418*f*
Electro codesign, 343
Electro-absorption (EA), 70, 249
  modulation mechanisms, 83–85
  modulators, 93–94
Electro-absorption modulators (EAMs), 73–76, 78–80, 83–87, 85*f*, 86*f*, 88*t*, 89*t*, 243. *See also* Group IV modulators; III–V modulators
  Franz–Keldysh effect, 73–75
  nonlinear optical effects, 76–77
  quantum-confined stark effect, 75–76
Electro-optic effect (EO effect), 70
  E-O switching, 274–276, 280
  EO cladding materials, 101
  EO Mach–Zehnder interferometer switch fabric, 280–281
  EO modulators, 159–161
Electro-refraction effects, 70
Electro-static discharge testing (ESD testing), 444
Electromagnetic (EM)
  simulations, 343
  solvers, 339
Electronic and photonic cointegrated circuits (EPIC), 139–140
Electronic chips, 335
Electronic circuits, 345
Electronic connectivity, 348–349
Electronic design automation (EDA), 371–372
Electronic dispersion, 461
Electronic packet switches (EPS), 266
Electronic photonic design automation (EPDA), 345, 421–422
Electronic–photonic codesign (E/O codesign), 346
Electronics integration, reliability of, 450–452, 450*f*, 451*f*
ELS. *See* External laser source (ELS)
EM. *See* Electromagnetic (EM)
EME. *See* Eigen mode expansion (EME)
EMLs. *See* Externally modulated lasers (EMLs)
Energy requirements, 54
Enterprise data center (EDC), 4, 11
Enterprise data center (EDC), 1–8
EO effect. *See* Electro-optic effect (EO effect)
EoM. *See* Error of Measurement (EoM)
EPDA. *See* Electronic photonic design automation (EPDA)
EPIC. *See* Electronic and photonic cointegrated circuits (EPIC)
Epitaxial BTO layers, 106–107
Epitaxy, 132–136
EPS. *See* Electronic packet switches (EPS)
ER. *See* Extinction ratio (ER)
Error compensation schemes, 234–235
Error of Measurement (EoM), 441
ESD testing. *See* Electro-static discharge testing (ESD testing)
Etching, 380, 394–395
Ethernet, 4, 7–8
Evanescent couplers, 174–175
Explicit Congestion Notification (ECN), 16
External cavities, 43
External laser source (ELS), 256–257
External modulation scheme, 49
Externally modulated lasers (EMLs), 233–234, 248–250
Extinction ratio (ER), 70, 239

F

F16 fabric, 14–15, 14*f*
FA. *See* Failure analysis (FA)
Fabric architecture, 322–324
Fabrication process, 48, 103, 174, 295–296, 338, 462–463
Fabry–Pérot lasers (FP lasers), 36, 40, 56
Facebook data centers, 266–267
Failure analysis (FA), 447
Failure mode effects analysis (FMEA), 444–445
Failures in Time (FIT), 29–30, 440, 449
"Fair-share" scheduler, 6–7
FAs. *See* Fiber arrays (FAs)

Fat Tree (FT) networks, 5, 8
FAU. *See* Fiber array units (FAU)
FCB assembly process, 286–287
FCPD effect. *See* Free-carrier plasma dispersion effect (FCPD effect)
FCPD-based optical modulators, 73
FDL. *See* Fiber delay line (FDL)
FDTD. *See* Finite difference time domain (FDTD)
FEC. *See* Forward error correction (FEC)
Feedback tolerance, 53
Feeding circuit level design flow, 343–345
FEOL. *See* Front-end-of-line (FEOL)
Fiber
  coupling to silicon photonic waveguides, 415*f*
  interfaces, 16
    to silicon photonics PICs, 457*f*
  reliability of fiber interfaces to PICs, 455–457
Fiber array units (FAU), 455–456
Fiber arrays (FAs), 385–386, 412–413
Fiber delay line (FDL), 296
Fiber to the Building/Home (FTTx), 142
Fiber-pigtail solutions, 52
Fiber-to-chip couplers, 172–175, 172*t*
  edge couplers, 173
  evanescent couplers, 174–175
  grating couplers, 174
Fiber-to-Si PIC coupling mechanism, 52
Figures of merit (FoM), 69
FILL, 363–364
Filters, 175–185, 372–373
Fine wavelength tuning, 39
Finite difference eigenmode solvers, 372
Finite difference time domain (FDTD), 342, 372
Finite impulse response (FIR), 176
FIR. *See* Finite impulse response (FIR)
FIT. *See* Failures in Time (FIT)
FK type modulators. *See* Franz–Keldysh type modulators (FK type modulators)
FKE. *See* Franz–Keldysh effect (FKE)
Flat butterfly, 9–10
Flat-top spectral response, 179
Flip-chip process, 417
  bonded lasers, 43
  bonding
    assembly process, 286–287
    gain integration with, 285
FMEA. *See* Failure mode effects analysis (FMEA)
FoM. *See* Figures of merit (FoM)
Footprint, 70
Formalisms, 345–346
Forward error correction (FEC), 21, 234–235
Forward looking heterogenous system designs, 309–310
FP lasers. *See* Fabry–Pérot lasers (FP lasers)
FPGAs, 27–28
FPP Modules. *See* Front-panel pluggable modules (FPP Modules)
Franz–Keldysh effect (FKE), 70, 73–75, 74*f*, 83–85
Franz–Keldysh type modulators (FK type modulators), 243
Free space micro optics (FSMO), 428–429, 433
Free spectral range (FSR), 177
Free-carrier plasma dispersion effect (FCPD effect), 71–73
  electro-absorption modulation, 73–76
  interferometric and resonant devices employing, 72*f*, 74*f*, 82–83, 84*t*
Front-end-of-line (FEOL), 107, 379
  approach, 140–141
  integration, 140
Front-panel pluggable modules (FPP Modules), 19, 21
Front-plate pluggable optical modules, 20*t*, 254
FSMO. *See* Free space micro optics (FSMO)
FSR. *See* Free spectral range (FSR)
FT. *See* Fat Tree (FT)
FTTx. *See* Fiber to the Building/Home (FTTx)
Fully monolithic photonic integration platform, 406–407

G

400-G coherent ecosystem, 209–210
100-G CWDM4 transceivers, 250
400-G DR4 transceivers, 250
400-G FR4 transceivers, 250
64-G integrated coherent receiver photonic-integrated circuit responsivity, 146*f*

100-G LR4, 200-G FR4 transceivers, 250
GaAs. *See* Gallium arsenide (GaAs)
Gain bandwidth product (GBP), 142
Gain integration with flip-chip bonding, 285, 286*f*
Gallium arsenide (GaAs), 35
Gaussian beam, 416–417
10-Gb/s APD, 142–143
53-Gb/s NRZ sensitivity testing, 145
GbE. *See* Gigabit Ethernet (GbE)
GBP. *See* Gain bandwidth product (GBP)
GDSII, 344
Ge/Si APDs. *See* Ge/Si avalanche photodiodes (Ge/Si APDs)
Ge/Si avalanche photodiodes (Ge/Si APDs), 128
Germanium (Ge), 73–75, 244, 396–397
  advanced integration of Ge devices with CMOS electronics, 139–142
  dislocations and lattice strain in, 132–136
    epitaxial growth of Ge on Si by CVD, 132
    tensile lattice strain by thermal stress, 134–136
    threading dislocation density in Ge, 133–134
  Ge PDs, 140–141
    processing challenges with, 137–139
  Ge p–i–n photodetectors, 137*f*
  Ge/Si APDs, 147
  Ge/Si heterojunction APDs meet current system, 148
  Ge/Si heterojunction avalanche photodiodes, 142–143
  $Ge_2Sb_2Te_5$-on-silicon nonvolatile switch, 287–288
  high-speed photodetectors, 129–131
    Ge/Si *vs*. InGaAs/InP high-speed p–i–n photodiodes and APDs, 130*f*
  thin films, 136
  TWPDs, 128
GF. *See* GlobalFoundries (GF)
Gigabit Ethernet (GbE), 53–54
Glass transition temperature (Tg), 456
GlobalFoundries (GF), 60, 391–392
Google, 316
Graphene
  graphene-based modulation, 102
  optical absorption of DLG on SOI passive waveguide, 103*f*
  on Si, 102–103
Grating couplers, 172–174
Grating-based single-wavelength laser, 38
Group delay, 87–90
Group IV modulators, 80–93. *See also* Electro-absorption modulators (EAMs); III–V modulators
  electro-absorption modulators, 83–87
  interferometric and resonant devices, 82–83
  modulation methods, 80–82
  modulator driver amplifiers, 87–93

H

HBTs. *See* Heterojunction bipolar transistors (HBTs)
HCS. *See* Heterogeneous computing systems (HCS)
Heat conduction, 440
HEMTs. *See* High electron mobility transistors (HEMTs)
Heteroepitaxially-grown monolithic lasers, 48–49
Heterogeneous computing systems (HCS), 303, 324
  applications space with requirements, 305–308
  cost, 320–322
  fabric architecture, 322–324
  interfaces, 308–314, 309*f*, 310*f*, 310*t*, 311*t*, 313*f*, 314*f*
  photonic-integrated circuits for heterogeneous computing systems, 324–329
  power and cooling, 318–319
  problem statement, 303–305
  system level reliability, resiliency, BER, lifetime, 314–317
Heterogeneous integrate III–V material, 253
Heterogeneous integrated DFB laser arrays, 60–61
Heterogeneous integration techniques, 43–45, 51, 53, 60–62, 219, 420
  approaches to achieve narrow line width heterogeneous III–V/Si lasers, 61*f*
  of III–V materials, 46

Intel's non hermetic packaged heterogeneous InP/Si transceiver picture and configuration, 44*f*
reliability of lasers integrated into PICs through, 453–454
on silicon, 97–108
$BaTiO_3$ on Si, 105–108
graphene on Si, 102–103
III–V on silicon, 97–99
lithium niobate on Si, 104–105
organic electro-optic materials on Si, 99–102
Heterogeneous laser integration approach, 389–390
Heterogeneous photonic integrated switch fabrics, 285–288
gain integration with flip-chip bonding, 285
$Ge_2Sb_2Te_5$-on-silicon nonvolatile switch, 287–288
silicon-semiconductor optical amplifier gate wavelength selective switch, 286–287
Heterogeneous photonics integration platform, 386–391
design considerations, 387–388
hardware performance, 389–390
heterogeneous integration, 390–391
integration approach, 388–389
Heterogeneous-integrated lasers, 49, 60, 63–64
Heterogeneously integrated ring lasers, 37–38
Heterojunction bipolar transistors (HBTs), 399–400
Hewlett Packard Enterprise, 27
HGRID, 13, 15*f*
Hierarchical design methodologies, process design kits and use of, 346
Hierarchical topologies, 6–7
High capacity switching networks, 267–268
High density Si waveguides, 161–163
High electron mobility transistors (HEMTs), 399–400
High quality PDKs, 355
High radix switches, 5–6
High refractive index contrast, 358
of silicon substrate, 173
High temperature operating life (HTOL), 443
High temperature storage (HTS), 443
High volume manufacturing (HVM), 369, 393–394, 465
High-bandwidth switches, 254
High-coherence single longitudinal mode laser, 56
High-index contrast Si waveguides, 140–141
High-order harmonic mode-locking scheme, 40
High-performance computing (HPC), 1–8, 265, 439
considerations for schedulers, 6–7, 6*f*
impact of emerging artificial intelligence workloads, 3–4
historical trends in, 7–8, 9*f*
link technologies, 4–5
optical interconnect architectures based on optical switches, 288–294
reliability standards for DC and HPC applications, 443–444
topologies, 5–6
workload requirements, 2–3, 2*f*
High-Performance Microelectronics (IHP), 391–392
High-performance photonic devices, 140–141
High-performance silicon photonic devices, 35
High-power disaggregated laser sources, 63
High-power lasers, 55
High-Q ring resonators, 43
High-quality factor microring resonators, 372
High-responsively PDs. *See* High-responsively Photodiodes (High-responsively PDs)
High-responsively Photodiodes (High-responsively PDs), 55–56
High-sensitivity photodetection, 128
High-speed Ge/Si lateral p–i–n photodiodes, 126–127
High-speed photodetectors
advanced Ge photodiode processing on CMOS platform, 136–142

High-speed photodetectors (*Continued*)
avalanche photodetectors, 142–149
epitaxy, 132–136
photodetectors in data center architecture, 123–132
device structures, 126–129
Ge *vs*. III–V high-speed photodetectors, 129–131
performance parameters and waveguide–detector integration paradigm, 124–126
TxRx subsystem requirements, 131–132
photodetectors in data centers, 149
High-temperature cofired ceramic (HTCC), 457–458
High-yield silicon photonic devices, 35
Higher bandwidth link technologies, 5
Highly scalable fast optical circuit switch networks, 292–294, 293*f*, 294*f*
HIPPI6400, 5
HPC. *See* High-performance computing (HPC)
HTCC. *See* High-temperature cofired ceramic (HTCC)
HTOL. *See* High temperature operating life (HTOL)
HTS. *See* High temperature storage (HTS)
HVM. *See* High volume manufacturing (HVM)
Hybrid integrated lasers, 48
Hybrid integration, 42–43, 58–60, 378–379
Hybrid laser integration, 60
Hybrid laser-to-PIC packaging, 420
Hybrid photonics integration platform, 378–386
design considerations, 378–382, 382*f*
hardware performance, 382–383, 383*f*, 384*f*
hybrid integration, 384–386, 385*f*, 386*f*
Hyperconvergence, 9–10
Hyperscale data centers, 2–3
DC showing hierarchy of interconnects used and connectivity of multiple DCs, 13*f*
Web 2.0, 13–19
data center region showing six data halls connected with HGRID, 15*f*
F16 fabric logical representation, 14*f*
Meta DC physical layout—two data halls, 14*f*
HyperX, 3, 9–10

I

I/O devices. *See* Input/output devices (I/O devices)
IA. *See* Implementation agreement (IA)
IB. *See* InfiniBand (IB)
ICR. *See* Integrated coherent receiver (ICR)
ICs. *See* Integrated circuits (ICs)
IEEE 802.3 Ethernet working group, 234–235
IHP. *See* High-Performance Microelectronics (IHP)
III–V APDs, 147–148
III–V epi wafer, 46
III–V high-speed photodetectors, 129–131
III–V modulators, 93–97. *See also* Electro-absorption modulators (EAMs); Group IV modulators
conventional III–V modulators, 94–96
membrane III–V modulators, 96–97
vertical-cavity surface-emitting lasers, 95*t*
III–V on silicon, 97–99, 98*f*, 99*f*
IILD. *See* Impurity-induced layer disordering (IILD)
IIR. *See* Infinite impulse response (IIR)
IIR optical filters, 176–177
IL. *See* Insertion loss (IL)
ILD. *See* Interlayer dielectric (ILD)
IM. *See* Intensity modulation (IM)
IM and direct detection (IM/DD), 101
IM/DD. *See* IM and direct detection (IM/DD)
IMDD. *See* Intensity modulation and direct detection (IMDD)
Implementation agreement (IA), 221, 226–227
Impurity-induced layer disordering (IILD), 401–402
Indium gallium aluminum arsenide (InGaAlAs), 57, 369
Indium gallium arsenide phosphide (InGaAsP), 369
Indium phosphide (InP), 35, 274–276, 369
coherent modulator with integrated optical gain, 212*f*

InP III−V photonics integration platform, 399−406
challenges, 404−406
design considerations, 400−401
InP fabrication approach, 401−403
InP III−V photonics, 403−404
integration with SiGe electronic-integrated circuits, 403
InP optical switch fabric, 278−280
active−passive semiconductor optical amplifier space and wavelength switch, 282−285
active−passive semiconductor optical amplifier space switch, 278−279
all-active semiconductor optical amplifier space switch, 278
InP-based lasers, 57−58
InP's QCSE technology, 217
modulators, 212
photonic switch, 274
wafers, 51
Inductive peaking, 92−93
Industry photonics suppliers, 54
InfiniBand (IB), 7−8
Infinite impulse response (IIR), 176
Inflection points and challenges, 24−30
interconnect cost targets, 28−29
power and thermal, 25, 26*f*, 27*f*
reliability, 29−30
resource disaggregation, 25−28
InGaAlAs. *See* Indium gallium aluminum arsenide (InGaAlAs)
InGaAsP. *See* Indium gallium arsenide phosphide (InGaAsP)
InP. *See* Indium phosphide (InP)
Input/output devices (I/O devices), 244−245, 423
I/O edge coupler integration, 394−395
Insertion loss (IL), 70, 102−103
Integrated circuits (ICs), 255, 338
Integrated coherent optics, progression of, 213*f*
Integrated coherent receiver (ICR), 146
Integrated coherent transmitter and receiver, 146
Integrated device manufacturers, 315
Integrated lasers, 58−63, 390
heterogeneous, 60−62
hybrid, 58−60
integration issues, 46−53
laser solutions, 41−46
monolithic, 62−63, 63*f*
outlook, 63−65
next 5 years, 63−64, 64*f*
next 20 years, 65
performance requirements, 53−58
energy requirements, 54
intensity noise, 57
operation temperature, 57−58
output power, 55−56
spectral characteristics, 56−57, 56*f*
threshold and slope efficiency, 55
semiconductor laser line width progression, 59*f*
state-of-the-art, 58−63
disaggregated, 58
integrated, 58−63
types of sources needed for datacom, 35−41
Integrated optical switch technologies, overview of, 274−277, 275*t*
Integrated photonic building blocks for data center interconnect, 241−245
laser integration, 242−243
modulators, 243−244
passives, 244−245
receiver, 244
Integrated photonic(s), 35, 346
components, 35
devices, 344
waveguides, 160−161
Integrated PIC industry, 365
Integrated sources, 42−46
Integration process, 107
of BTO-based photonic modulators, 108*f*
future directions in, 186−187
issues, 46−53
cost, 51
coupling to Si waveguide, 51−53
heterogeneous InP/Si lasers and monolithic GaAs-on-Si lasers, 50*f*
optical stability and feedback tolerance, 53
process integration, 46−47
process temperatures, 48
reliability requirements, 49−51
yield, 48−49
to support scaling to 400-to 800-G transceivers, 215*f*

Intel approach, 60–61
Intel Foundry Services, 60
Intensity modulation (IM), 78–80
Intensity modulation and direct detection (IMDD), 201, 234–235
  to coherent, evolution of DCI link technologies from, 206–210
Intensity noise, 57
Intentional surface doping, 139
Interconnect figures of merit, 23–24
  interconnect efficiency *vs*. bandwidth, 24*f*
  interconnect efficiency *vs*. reach, 23*f*
Interconnect technology, 9*f*, 23
Interdatacenter interconnects, 225
Interferometers, 168
  filters, 177–179, 180*t*
Interferometric approach, 81–82
Interferometric devices, 78–80
  employing free-carrier plasma dispersion effect, 82–83
Interferometric photonic modulators, 81, 81*f*
Interferometry, 170–171
Interlayer dielectric (ILD), 140–141
International Roadmap for Device and Systems, The, 305–307
Intradata center interconnects, 225–226
  pluggable transceivers for, 245–253
Intradata center link requirements and standards, 234–237
  data center bandwidth and transceiver growth, 234*f*
  for different QSFP generations, 235*t*
  switch chassis design configurations, 236*f*
Inverse design, 164, 169
Ionization coefficient, 143
IPoDWDM shift, market impact of, 223–224, 224*f*

K

Keep-out zone (KOZ), 381
Kerr comb lasers, 40–41
Kerr effect, 70, 76–77
Kerr nonlinear frequency combs, 39–40
Kirchhoff's conservation laws, 336
KOZ. *See* Keep-out zone (KOZ)

L

LaMP. *See* Laser micropackage (LaMP)
Land grid array (LGA), 457–458
Laser micropackage (LaMP), 452–453, 453*f*
Laser out-coupling, 52
Laser(s), 53, 274
  diodes, 452
  disaggregated sources, 41–42
  integrated sources, 42–46
    heterogeneous integration, 43–45
    hybrid integration, 42–43
    monolithic integration, 45–46
  integration, 242–243, 395
    laser integration schemes, 243*f*
  Lorentzian linewidth, 57
  output power, 55–56
    requirement, 55
  solutions, 41–46
Laser soliton microcomb devices, 60
Latency, 3, 314*f*
Lateral p–i–n PD structure, 138
Layout-dependent variation models, 361
Layout *vs*. schematic (LVS), 371–372
Lead zirconate titanate, 243–244
LER. *See* Line edge roughness (LER)
4-level pulse amplitude modulation (PAM4), 234–235
LGA. *See* Land grid array (LGA)
LiDAR, 142
Light, 81, 99–101
  reliability of light source, 452–453, 453*f*
$LiNbO_3$'s Pockels effect, 217
Line edge roughness (LER), 359
LINPACK benchmark, 7
Liquid cooling, 25, 319
Lithium niobate (LN), 76
  on Si, 104–105, 104*f*
Lithography, 374–376
LN. *See* Lithium niobate (LN)
LN EO modulators, 104
Long reach (LR) multimode transceivers, 233–234
Long-haul fiber optic communications, 411–412
Lower radix topologies, 5
Lumped elements, 90–91
LVS. *See* Layout *vs*. schematic (LVS)

M

Machine, The, 27
Machine learning, 288
  hardware, 18–19

meta Zion machine learning platform, 18*f*
Zion connectivity and interconnect fabrics, 18*f*
model training, 324
Mach−Zehnder interferometers (MZIs), 71, 78−80, 243−244, 269, 372
modulators, 82
Mux/Demux demonstrations, 180*t*
Mach−Zehnder modulators (MZMs), 90−91, 94, 179−181, 208, 243−244, 390, 441, 464
Magnesium oxide (MgO), 105−106
Manhattanization of curvilinear structures, 374−376
Manufacturability, design for, 358−364
Manufacturing FMEA (MFMEA), 445
Manufacturing process, 160−161, 347−348, 352−353
Manufacturing variations, 358−359
manufacturing variation-aware photonic models, 359−361
Marvell Silicon Photonics modules, 459
Mask process correction (MPC), 374−376
Mask rule check (MRC), 373−374
Mature semiconductor lasers, 35
Maxwell's equations, 337, 339−342
MBE. *See* Molecular-beam epitaxy (MBE)
Mean time between failure (MTBF), 449
Mean Time to Failure, 29−30
Membrane III−V modulators, 96−97
Memory disaggregation, 29
Memory-intensive workloads, 18−19
MEMS, 274−276
MEMS-based switch, 274
Mesh networks, 5
Meta DC physical layout—two data halls, 14*f*
Meta platforms, 13
Meta switch and optics deployment roadmap, 17*t*
Meta Zion machine learning platform, 18*f*
Metal lines, 391
Metal organic chemical vapor deposition (MOCVD), 46
Metal−semiconductor−metal (MSM), 127
MFD. *See* Mode-field diameter (MFD)
MFMEA. *See* Manufacturing FMEA (MFMEA)
Micro Lens Arrays (MLAs), 416
Micro-optics for integrated photonics, 3D printing of, 427−434
Microbumps, 381−382
Microlens on photonic device, 416*f*
Microresonator ring (MRR), 327
Microring modulators (MRMs), 38, 244
Microring resonant, 176−177
Microring resonator (MR), 179−181, 274−276
Microring WDM Mux/Demux demonstrations, 182*t*
Microring-based WDM Mux/Demux interconnect, block diagram of, 181*f*
Microsoft Research, 323−324
Microtransfer printing method, 44
Mid-infrared transmission (MIR), 163
Minipack (12.8 Tb/s modular design), 14−15, 15*f*
MIR. *See* Mid-infrared transmission (MIR)
MLAs. *See* Micro Lens Arrays (MLAs)
MM optical cable. *See* Multimode optical cable (MM optical cable)
MMI. *See* Multimode interferometer (MMI)
MOCVD. *See* Metal organic chemical vapor deposition (MOCVD)
Mode evolution-based rotators, 171
Mode solvers, 339−340, 340*f*, 341*f*, 372
Mode-coupling-based devices, 171
Mode-field diameter (MFD), 414−415
Mode-locked comb lasers, 39−40
Mode-locked laser, 40
Mode-locked semiconductor laser, 40
Model parallelism, 316
Modeling, notes on electrical *vs*. optical simulation and, 353−355
Modern computing workloads, 303
Modulation mechanisms, 70−77
free-carrier plasma dispersion effect, 71−73
thermo-optic effect, 71
Modulation methods, 80−82
interferometric photonic modulators, 81
interferometric vs. resonant approaches, 82, 83*t*
resonant photonic modulators, 81−82
Modulators, 214−215, 243−244, 344
driver amplifiers, 87−93
bandwidth and group delay, 87−90
circuit topology, 90−92

Modulators (*Continued*)
codesign of driver and modulator, 92−93
voltage swing and power efficiency, 90
Module level reliability, 459−462
Molecular-beam epitaxy (MBE), 46
Monolithic electro-optical processes, process design kits for, 356−357
Monolithic electronic−photonic integration, 123
Monolithic integration, 45−46, 45*f*, 62−63, 65
Monolithic photonics integration platform, 391−399
design considerations, 392−393
hardware performance, 396−398
45SPCLO CMOS technology, 396*t*
CMOS technology comparison, 397*t*
GF45SPCLO technology, 398*t*
I/O edge coupler integration—V-groove, 394−395
integration approaches, 393−394
laser integration, 395
monolithic integration, 398−399
Monolithic silicon photonics technology, 398−399
Monolithic tunable optical router (MOTOR), 282−283
Moore's law, 7
MOSCAP design, 97−99
MOTOR. *See* Monolithic tunable optical router (MOTOR)
MPC. *See* Mask process correction (MPC)
MQW. *See* Multiquantum well (MQW)
MR. *See* Microring resonator (MR)
MRC. *See* Mask rule check (MRC)
MRMs. *See* Microring modulators (MRMs)
MRR. *See* Microresonator ring (MRR)
MSA. *See* Multisource agreement (MSA)
MSM. *See* Metal−semiconductor−metal (MSM)
MTBF. *See* Mean time between failure (MTBF)
Multi wavelength lasers, 36−38
FP, 36
ring, 37−38
Multichip module. *See* 2D packaging
Multifiber connectors, 21
Multimicrons MM waveguides, 244−245
Multimode interferometer (MMI), 167, 169
Multimode optical cable (MM optical cable), 233
Multimode technologies, 24
Multiorder MZI filters, 179
Multiple optical vendors, 315
Multiplexers, 175−185
Multiquantum well (MQW), 75−76, 400
Multisource agreement (MSA), 219−220, 312
Multistage architectures, 281
Multiwavelength lasers, 36, 56, 60, 63
Multiwavelength sources, 61
Mux. *See* Optical multiplexer (Mux)
Mux/Demux, 177−179
64 × 64 MZI-based thermo-optical switch, 278−279
MZIs. *See* Mach−Zehnder interferometers (MZIs)
MZMs. *See* Mach−Zehnder modulators (MZMs)

N

N-stage planar network, 273
Nanoseconds optical switching, disaggregated architecture based on, 291−292
Nanoseconds SOA-based optical switches (NOS), 291−292
Near infrared (NIR), 163
Network control plane, 288−290
Network hardware, 14−15, 15*f*
Network interface cards (NICs), 233
Network interface controllers (NICs), 18−19
Network-assisted signaling, 16
New Product Introduction (NPI), 444
Next-generation network switches, 29
NICs. *See* Network interface cards (NICs); Network interface controllers (NICs)
NIR. *See* Near infrared (NIR)
28nm FD-SOI process, 90−91
Nonlinear optical effects, 76−77
Kerr effect, 76−77
Pockels effect, 76
Nonreturn to zero (NRZ), 234−235
Nonvolatile materials, 285
NPI. *See* New Product Introduction (NPI)
NRZ. *See* Nonreturn to zero (NRZ)
NVLink/NVSwitch, 3

O
OBO. *See* On-board optics (OBO)
OCAP. *See* Out-of-control action plans (OCAP)
OCS. *See* Optical-circuit switching (OCS)
OCSs. *See* Optical circuit switches (OCSs)
OFCs. *See* Optical frequency combs (OFCs)
Off-package light source, 420
OIF. *See* Optical internetworking forum (OIF)
OIF 400ZR specifications, 221
OMA. *See* Optical modulation amplitude (OMA)
On-board optics (OBO), 236–237, 309
On-chip
  routing devices, 244–245
  transmission line, 92
On-going process, 440
One-time poling process, 99–101
Ongoing reliability monitoring test (ORM), 465
On–off keying (OOK), 203–204
  transmitters, 439
OOK. *See* On–off keying (OOK)
OPC. *See* Optical proximity correction (OPC)
Open-drain amplifier, 91–92
OpenROADM MSA, 221
OpenROADM specification, 223
OpenZR + , 223
Operation temperature, 57–58
OPLLs. *See* Optical phase locked loops (OPLLs)
OPS. *See* Optical-packet switching (OPS)
OPSquare architecture, 287
Optical absorption, 136
Optical and electrical connections (I/Os), 411–412
Optical bandwidth, 70
Optical bonds, 455–456
Optical cables, 28, 123
Optical circuit switches (OCSs), 288–290
Optical codesign, 343
Optical communication, 242–243
Optical community, 29
Optical coupling technique, 125, 417
Optical design optimizations, 106–107
Optical fiber, 335
  coupling, 172–173
Optical frequency combs (OFCs), 40
Optical I/O PICs, 325–326
Optical I/O pins, 140
Optical integration, levels of, 214
Optical interconnect technologies, 266
Optical interconnection, 140
Optical interconnects, 1, 23–24, 201
Optical interfaces, 16
  design considerations, 312
Optical internetworking forum (OIF), 21–22
Optical interposers, 417
Optical lattice filters, 179
Optical microscopy, 145*f*, 447
Optical modulation, 70, 240
Optical modulation amplitude (OMA), 239
Optical modulators
  group IV modulators, 80–93
  heterogeneous integration on silicon, 97–108
  historical review, 77–80, 78*f*
  III–V modulators, 93–97
  modulation mechanisms, 70–77
Optical modules, 23
  approaches, 313*f*
  form factors, 19–23, 20*t*
Optical multiplexer (Mux), 175
Optical Mux/Demux, 176
Optical networks, 142
  case studies, 1–19
    enterprise data center, 1–8
    high-performance computing, 1–8
    hyperscale data centers, 2–3
Optical packaging process, 414–417
  fiber coupling to silicon photonic waveguides, 415*f*
  microlens on photonic device, 416*f*
  photonic devices, 415*f*
  ZEMAX simulations, 416*f*
Optical packaging process, 416
Optical phase locked loops (OPLLs), 226
Optical proximity correction (OPC), 167, 363–364, 374–376
Optical signal-to-noise ratio (OSNR), 269
Optical simulation, notes on, 353–355
Optical stability, 53
Optical switch architectures, 270
Optical switch fabric architectures, 270–273
  banyan-type blocking networks, 270–271
  figures of merit for, 272*t*

Optical switch fabric architectures (*Continued*)
rearrangeable nonblocking networks, 271–273
schematic representation of, 271*f*
wide-sense and strictly nonblocking networks, 273
Optical switches, 168, 270
Optical switching
distributed deep learning training based on, 288–290
improved resource utilization through bandwidth steering enabled by, 290
key figures of merit for optical switching networks, 266–270
device metrics, 268–270
system metrics, 266–268
Optical transceiver, 245–246, 255, 414
Optical transmission system, 403
Optical waveguides, 350
Optical-circuit switching (OCS), 322
Optical-packet switching (OPS), 323
Optics, 29
Optoelectronic device, 443–444
Optoelectronic test, 376
Organic electro-optic materials on Si, 99–102, 100*f*
Organic substrates, 309
OSAT. *See* Outsourced Semiconductor Assembly and Test (OSAT)
OSNR. *See* Optical signal-to-noise ratio (OSNR)
Out-of-control action plans (OCAP), 405
Output power, 55–56
Outsourced Semiconductor Assembly and Test (OSAT), 450

P

Package substrate technologies, 310*f*
p-cell. *See* Parameterized cells (p-cell)
Package-by-package process, 411–412
Package-level process, 411–412
Packaging, 352–353
3D printing of micro-optics for integrated photonics, 427–434
configurations for PICs and ASICs, 214*f*
design rules and standardization, 420–421
electrical packaging, 417–420
hybrid laser-to-PIC packaging, 420
optical packaging, 414–417
packaging design rules and standardization, 420–421
PIC packaging technologies, 412, 412*f*
software design tools, 421
testing, 421–427
PIC testing equipment, 423–425
roadmap for PIC testing, 426–427
standardized PIC testing protocols, 425–426
transceiver package design, 412–414
types of, 214
Packaging/assembly design kits (ADKs), 420–421
PAM. *See* Pulse amplitude modulation (PAM)
PAM4 modulation, 239
Parallel Computer Interface (PCIe Gen4), 25–26
Parallel computing, 5
Parallel single mode (PSM), 24, 42, 246–248
Parameterized building blocks, 348–349
Parameterized cells (p-cell), 372–373
Passive devices, 92, 159, 344
Passive optical components, 244–245
Passive optical devices
fiber-to-chip couplers, 172–175
polarization diversity, 170–171
routing elements, 164–168
splitters/combiners/couplers, 168–170, 168*t*
adiabatic, 170
branch, 169
directional, 169–170
multimode interferometers, 169
survey of key, 160–175
waveguides, 160–164
Passive silicon photonic devices, 160*f*
future directions, 185–187
future directions in design and testing, 186
future directions in integration, 186–187
multiplexers/demultiplexers and filters, 175–185, 178*t*
arrayed waveguide grating and echelle diffraction gratings Mux/Demux, 183–185

interferometer filters and Mux/Demux, 177–179
requirements and current trends, 175–177
resonator Mux/Demux, 179–181
waveguide Bragg grating Mux/Demux and filters, 181–183
survey of key passive optical devices, 160–175
Passives, 107, 244–245
Path-independent loss (PILOSS), 270, 273
32 × 32 path-independent loss switch, 278
PBS. *See* Polarization beam splitters (PBS)
PCells, 348–349
PD. *See* Photodetectors (PD); Photodiode (PD)
PDKs. *See* Process design kits (PDKs)
Peaking techniques, 217
Performance metrics, 70
Performance requirements, 53–58
PFMEA. *See* Process FMEA (PFMEA)
Phase modulation, 81
Phosphorus (P), 137
Photochemical degradation, 456–457
Photodetectors (PD), 123–124, 159, 239, 377
building block, 400–401
in data centers, 149
Photodiode (PD), 203–204, 446, 458–459, 464
Photoluminescence (PL), 48
Photon absorption, 125
Photon-photon resonance effect, 249
Photonic building blocks, 397–398
Photonic chips, 335, 357
Photonic circuits, 167–168, 338, 345
Photonic connectivity, 348–349
Photonic devices, 140, 339–340, 417
co-packaging of, 419*f*
design, 338–345
electro/optical codesign, 343
feeding circuit level design flow, 343–345
mode solvers, 339–340
propagation solvers, 340–342
technology computer aided design, 338–339
Photonic *I/O*s, Photonic-integrated circuits for, 324–327
Photonic integration, 226, 439
photonic integration/copackaging for data center interconnects, 254–257
Photonic interconnect fabrics, 304–305
Photonic modulators, 179–181
Photonic packaging, 420–421
Photonic PDKs, 356
Photonic subsystems, reliability of packaging for, 457–458
Photonic switch fabrics in data center/high performance computing networks
data center/high-performance computing optical interconnect architectures, 288–294
disaggregated architecture, 291–292
distributed deep learning training, 288–290
highly scalable fast optical circuit switch networks, 292–294
improved resource utilization through bandwidth steering enabled by optical switching, 290
heterogeneous photonic-integrated switch fabric, 285–288
InP optical switch fabric, 278–280
key figures of merit for optical switching networks, 266–270
overview of integrated optical switch technologies, 274–277
overview of optical switch fabric architectures, 270–273
remaining challenges, 294–296
silicon photonic switch fabric, 277–281
Photonic switch stage (PSS), 282–283
Photonic system designs, 364
Photonic TCAD of photonic device tools, 339
Photonic waveguide
connections, 349
layout, 350
Photonic wire bonds (PWBs), 60, 172–173, 433
Photonic-integrated circuits (PICs), 123–124, 233–234, 241, 343–344, 369, 411–412, 439
chips, 146
for coherent DCI, 210–219
coherent transceivers, 210–214
co-packaged optics, 214
directions for more integration, 218*f*

Photonic-integrated circuits (PICs) (*Continued*)
levels of electrical and optical integration, 214
performance details and limitations, 214–215
photonic-integrated circuit material platforms, 210
potential for further cost, size, and power reductions for coherent PICs, 215–219
types of packaging, 214
design, 345–357
back-annotation and postlayout simulation, 351
circuit simulation, 349–350
design flow, 347–348
design for manufacturability, design for test, and post-layout data prep, 352
formalisms, 345–346
future photonic-integrated circuit design flow requirements, 364–365
manufacturing, packaging, and test, 352–353
notes on electrical *vs*. optical simulation and modeling, 353–355
physical verification, 351–352
process design kits, 355–356
process design kits and use of hierarchical design methodologies, 346
process design kits for monolithic electro-optical processes, 356–357
schematic capture, 348–349
schematic driven layout, 350
design methodology, 345
for enabling disaggregated reconfigurable heterogeneous computing systems, 327–329, 329*f*, 330*f*
for enabling disaggregated reconfigurable heterogeneous computing systems, 327–329
fabrication schemes and methodology, 370–376
design rule check, 373
device design and simulation flow, 372
layout generation, 372–373
Manhattanization of curvilinear structures, 374–376
mask rule check, 373–374
physical design flow, 371–372
silicon photonics integration schemes, 370*f*
for heterogeneous computing systems, 324–329, 324*f*, 325*f*, 326*f*
photonic-integrated circuits for photonic *I/O*s, 324–327
in-line and end-of-line testing, 376–377, 377*f*
layouts, 164
loss bottlenecks and challenges, 239–241
manufacturing logistics, 371*f*
material platforms, 210
methodology for PIC reliability, 444–449
derisking, 445–448
design constraints, 448–449
failure mode effects analysis, 444–445
key steps for design for quality programs
packaging technologies, 412, 412*f*, 413*f*, 414*f*
photonic-integrated circuit in-line and end-of-line testing, 376–377
photonic-integrated circuit integration, fabrication, and manufacturing, 378–406
heterogeneous photonics integration platform, 386–391
hybrid photonics integration platform, 378–386
InP III–V photonics integration platform, 399–406
monolithic photonics integration platform, 391–399
photonic-integrated circuit outlook, 406–407
relevant industry trends, 369–370
reliability of DCI optical systems and field data, 458–462
module level reliability, 459–462
reliability requirements for DC/HPC application environment, 441–443, 442*t*
reliability standards for DC and HPC applications, 443–444
strategy for robust manufacturing quality, 462–465
high volume manufacturing, 463–464

on-going reliability monitoring, 464–465
wafer qualifications/certification, 463
subsystem and module assembly and integration, 449–458
burn-in, 452
reliability of fiber interfaces to PICs, 455–457
reliability of lasers integrated into PICs heterogeneous integration, 453–454
reliability of light source, 452–453
reliability of packaging for photonic subsystems, 457–458
reliability of photonics/electronics integration, 450–452
for switched network interconnects
data center interconnect, 233–241
data rate scaling vectors for future PIC generations, 259*f*
implementations of copackaged optics, 257*f*
integrated photonic building blocks, 241–245
photonic integration/copackaging, 254–257
photonic-integrated circuit scaling, 259–260
pluggable transceivers, 245–253
scaling datacenters to 102-Tbps nodes and beyond, 258–259
testing
equipment, 423–425
roadmap for, 426–427
standardized PIC testing protocols, 425–426
Photonic–electronic cosimulation, 354
Photonics
design for test and design for reliability, 364
design rule checking, 363–364
designer, 363–364
photonics-related applications, 35
reliability of photonics integration, 450–452
technologies, 245
Physical network latency, 16
Physical verification, 351–352
PICs. *See* Photonic-integrated circuits (PICs)
PILOSS. *See* Path-independent loss (PILOSS)
p–i–n, 128
PL. *See* Photoluminescence (PL)
Plasma dispersion effect, 70
Plasma etching, 363
Pluggable transceivers for intradata center interconnects, 245–253
block diagram of typical optical transceiver, 246*f*
comparing different types of 100 GBase transmitter specifications, 247*t*
directly modulated lasers, 248–249
externally modulated lasers, 249–250
Si photonics platform, 250–253
silicon PIC integration configurations, 246*f*
PMF. *See* Polarization maintaining fiber (PMF)
Pockel's effect, 70, 76, 211
Pol-mux. *See* Polarization multiplexing (Pol-mux)
Polarization beam splitters (PBS), 170–171
Polarization diversity, 170–171
polarization beam splitters, 170–171
polarization rotators, 171
scheme, 279
switch, 279–280
Polarization insensitivity, 279
Polarization maintaining fiber (PMF), 242
Polarization multiplexing (Pol-mux), 54
Polarization rotator-splitter, 255
Polarization rotators, 171
Polarization splitter/rotator (PSR), 397–398
Polarization-insensitive design, 280
Polarization-maintaining FAs, 424–425
Postlayout
data prep, 352
simulation, 351
Postmetallization annealing, 139
Proactive congestion avoidance methods, 3
Process, voltage and thermal variations (PVT variations), 362–363
Process design kits (PDKs), 337, 355–356, 371–372, 371*f*, 421–422
for monolithic electro-optical processes, 356–357
and use of hierarchical design methodologies, 346

Process FMEA (PFMEA), 445
Process integration, 46–47, 47*f*
Process technology, 347–348
Process temperatures, 48
Process variations, 44
Propagation solvers, 340–342
  BPM, 341–342
  EME and propagation, 341
  FDTD, 342
PSM. *See* Parallel single mode (PSM)
PSR. *See* Polarization splitter/rotator (PSR)
PSS. *See* Photonic switch stage (PSS)
Pulse amplitude modulation (PAM), 203–204
Pulsed vapor deposition (PVD), 380
Push–pull amplifier, 91
PVD. *See* Pulsed vapor deposition (PVD)
PVT variations. *See* Process, voltage and thermal variations (PVT variations)
PWBs. *See* Photonic wire bonds (PWBs)

Q
Q factor. *See* Quality factor (Q factor)
QAM. *See* Quadrature amplitude modulation (QAM)
QCSE. *See* Quantum-confined Stark effect (QCSE)
QD. *See* Quantum dot (QD)
QPSK signals. *See* Quadrature phase shift keying signals (QPSK signals)
QSFP. *See* Quad small form-factor pluggable (QSFP)
Quad small form-factor pluggable (QSFP), 19, 235, 235*t*, 412–413
Quadrature amplitude modulation (QAM), 93–94
  QAM16, 336
Quadrature phase shift keying signals (QPSK signals), 293–294
Quality factor (Q factor), 82
Quantum dot (QD), 40
  lasers, 57–58, 63*f*
  materials, 45–46, 53
  QD-based mode-locked lasers, 63
  single-section lasers, 40
Quantum well intermixing (QWI), 401–402
Quantum-confined Stark effect (QCSE), 70, 75–76, 75*f*, 212
  EAMs, 243
  QCSE-based devices, 85–87
QWI. *See* Quantum well intermixing (QWI)

R
Random early detection methods (RED methods), 3
RC. *See* Resistance–capacitance (RC)
RDLs. *See* Redistribution layers (RDLs)
Rearrangeable nonblocking
  networks, 271–273
  switch fabrics, 271
Receiver (RX), 244, 385
Receiver Optical Sub-Assembly (ROSA), 142–143
RED methods. *See* Random early detection methods (RED methods)
Redistribution layers (RDLs), 379, 418–419
Reflective terminations, 168
Relative intensity noise (RIN), 38, 248, 395
Reliability, photonics design for test and design for, 364
Reliability requirements, 49–51
Remote laser sources, 42
Remote terminals, 444
Residual gas analysis (RGA), 453
Resiliency, 314–317
Resistance–capacitance (RC), 124, 351
Resonant approaches, 82
Resonant devices employing free-carrier plasma dispersion effect, 82–83
Resonant modulators, 217
Resonant photonic modulators, 81–82, 82*f*
Resonators, 168
  Mux/Demux, 179–181, 181*f*, 182*t*
Resource disaggregation, 25–28
Resource utilization through bandwidth steering enabled by optical switching, improved, 290, 291*f*
RGA. *See* Residual gas analysis (RGA)
RIN. *See* Relative intensity noise (RIN)
Ring modulators (RM), 217, 379
Ring resonators, 90–91
  ring resonator-based modulator, 82*f*
  ring-resonator-based lasers, 37–38
Risk priority number (RPN), 444
RM. *See* Ring modulators (RM)

ROSA. *See* Receiver Optical Sub-Assembly (ROSA)
Routing, 371–372
  bends, 164
  elements, 164–168, 165*t*
  terminations, 167–168
  waveguide crossings and escalators, 164–167
RPN. *See* Risk priority number (RPN)
Ryu-based SDN controller. *See* Ryu-based software-defined networking controller (Ryu-based SDN controller)
Ryu-based software-defined networking controller (Ryu-based SDN controller), 288–290

S
SACM. *See* Separate absorption, charge, and multiplication (SACM)
Sagnac loops, 168
Salient fabrication, 440
Sampled-grating DBR lasers (SGDBR lasers), 39, 400
Scanning electron microscope (SEM), 132
Schedulers, considerations for, 6–7
Schematic-driven layout (SDL), 345–346, 350
SDL. *See* Schematic-driven layout (SDL)
SEG. *See* Selective epitaxial growth (SEG)
Segmented modulators, 92
Selective epitaxial growth (SEG), 132
Self-injection locking (SIL), 43
SEM. *See* Scanning electron microscope (SEM)
Semiconductor optical amplifier (SOA), 41, 268–269, 274, 286*f*, 386
Semiconductors, 73–75
  electronics, 422
  lasers, 51, 59*f*
Sensor markets, 142
Separate absorption, charge, and multiplication (SACM), 128
SerDes. *See* Serializer/deserialize (SerDes)
SERDES. *See* SERializer/DESerializer (SERDES)
Serializer/deserialize (SerDes), 256
SERializer/DESerializer (SERDES), 235–236, 238*t*
SGDBR lasers. *See* Sampled-grating DBR lasers (SGDBR lasers)
Shoreline density, 415–416
Short reach optical transceivers (SR optical transceivers), 233
Si_3N_4-based ring resonators, 59–60
Side-mode suppression ratio (SMSR), 248, 400
Signal-to-noise ratio (SNR), 234–235, 426–427
SIL. *See* Self-injection locking (SIL)
Silicon (Si), 102, 338, 392–394
  $BaTiO_3$ on, 105–108
  coupling to Si waveguide, 51–53
  graphene on, 102–103
  heterogeneous integration on, 97–108
  lithium niobate on, 104–105
  modulators, 213
  optical switches, 285
  organic electro-optic materials on, 99–102
  photonics platform, 250–253, 251*t*, 254*f*
  silicon-based PICs, 369
Silicon dioxide (SiO_2), 241
Silicon Franz–Keldysh modulator, recent progress in, 88*t*
Silicon germanium (SiGe), 378
Silicon nitride (SiN), 102, 241, 380, 388–389, 420
Silicon photonic (SiPh), 423
Silicon photonic integrated circuits (SiPICs), 41, 63, 233–234, 304–305
  passive platforms, 244–245
  products, 250
Silicon photonics (SiPh), 159–160, 164, 186–187, 274–276, 440, 450
  fabrication, 378
  interposer, 379–380
  modulators, 77–78
  process technologies, 387–388
  silicon photonic-integrated circuit integration configurations, 246*f*
  switch fabric, 274–281
    electro-optic Mach–Zehnder interferometer switch fabric, 280–281
    thermo-optic Mach–Zehnder interferometer switch fabric, 278–280
  technology, 249

Silicon photonics (SiPh) (*Continued*)
transceivers, 335
waveguides, 415*f*
Silicon-on-insulator (SOI), 77–78, 132, 159, 161, 241, 359
platform, 163, 378
wafers, 140
Silicon-organic hybrid (SOH), 99
EO modulators, 101
MZM concept, 99–101
Silicon-semiconductor optical amplifier gate wavelength selective switch, 286–287
measured fiber-to-fiber transmission and gain characteristics of semiconductor optical amplifier, 287*f*
Silicon–SOA hybrid wavelength selective switch, 285
Single frequency lasers, 38
DFB lasers, 38
distributed Bragg reflector lasers, 38
Single longitudinal lasers, 56
Single wavelength lasers, 36, 60
Single-layer graphene (SLG), 102
Single-mode fiber (SMF), 233–234, 414–415
Single-mode submicron waveguides, 244–245
SiP. *See* System-in-Package (SiP)
SiPh. *See* Silicon photonic (SiPh)
SiPICs. *See* Silicon photonic integrated circuits (SiPICs)
SLG. *See* Single-layer graphene (SLG)
SMF. *See* Single-mode fiber (SMF)
SMSR. *See* Side-mode suppression ratio (SMSR)
SNR. *See* Signal-to-noise ratio (SNR)
SOA. *See* Semiconductor optical amplifier (SOA)
SoC. *See* System-on-chip (SoC)
SOH. *See* Silicon-organic hybrid (SOH)
SOI. *See* Silicon-on-insulator (SOI)
Spatial variation models, 361
Spatially correlated process variations, 361–362
SPC. *See* Statistical process controls (SPC)
45SPCLO platform, 396–397
Split-carrier transmitters, 225–226
Splitters, 168–170, 372–373
Spot-size converter (SSC), 140
SR optical transceivers. *See* Short reach optical transceivers (SR optical transceivers)
SSC. *See* Spot-size converter (SSC)
Standard laser fabrication process, 46
Standard-based modules, 21–22
Statistical process controls (SPC), 405, 464
Stress and lifetime test (STRIFE test). *See* Derisking plan
Supercomputer, 290
Surface activation process, 48
Surface cleanness, 46
Switching element, 279
System design, PICs, 357–358
System level reliability, 314–317
System metrics, 266–268
System-in-Package (SiP), 424–425
System-on-chip (SoC), 325–326, 439

T

25T-100T network switches, scaling from, 215*f*
T-O switches. *See* Thermo-optic switches (T-O switches)
Tantalum (Ta), 380
51.2-Tbps Ethernet switch generation, 258
TCAD. *See* Technology computer-aided design (TCAD)
TDD. *See* Threading dislocation density (TDD)
TDEC. *See* Transmitter dispersion eye closure (TDEC)
TDECQ. *See* Transmitter dispersion eye closure quaternary (TDECQ)
TDs. *See* Threading dislocations (TDs)
TE polarization. *See* Transverse electric polarization (TE polarization)
TE0 modes. *See* Transverse-electric modes (TE0 modes)
Technology computer-aided design (TCAD), 338–339, 339*f*
TECs. *See* Thermo-electric coolers (TECs)
Telecom GR-468 Generic Reliability Assurance Requirements, 142–143
Telecommunication networks, 440
TEM. *See* Transmission electron microscope (TEM)
Tensile lattice strain by thermal stress, 134–136

Tensile strain, 130, 134–135
TeraPHY prototypes, 326
2-terminal modulators, 77–78
3-terminal modulators, 77–78
Terminations, 167–168
TFLN. *See* Thin-film lithium niobate (TFLN)
TGA. *See* Thermo Dynamic Analysis (TGA); Thermogravimetric analysis (TGA)
Thermo Dynamic Analysis (TGA), 456
Thermo-electric coolers (TECs), 242, 403, 457–458
Thermo-optic effect, 71
Thermo-optic Mach–Zehnder interferometer switch fabric, 278–280
  32×32 path-independent loss switch, 278
  64×64 Benes switch, 278–279
  polarization diversity switch, 279–280
Thermo-optic switches (T-O switches), 268–269
Thermogravimetric analysis (TGA), 456–457
Thin-film lithium niobate (TFLN), 104, 212*f*, 219, 243–244
Threading dislocation density (TDD), 48, 133–134
Threading dislocations (TDs), 133–134
3D printing of micro-optics for integrated photonics, 427–434
  application of, 428–433
  process for, 427–428
  roadmap for, 433–434
3D Torus networks, 5, 6*f*
Through silicon/substrate vias (TSVs), 370–371, 418–419, 451
Through-silicon vias (TSV), 379
TIA. *See* Transimpedance amplifier (TIA)
TIR. *See* Total internal reflection (TIR)
TM modes. *See* Transverse magnetic modes (TM modes)
TO-can. *See* Transistor-Outline-can (TO-can)
Tomahawk4 switches, 327
Top of rack (TOR), 12, 233, 266–267
TOR. *See* Top of rack (TOR)
TOSA. *See* Transmitter optical subassembly (TOSA)
Total internal reflection (TIR), 429
Tower Semiconductor Ltd., 60
Transceiver module architecture, 21–23, 22*f*
Transceiver package design, 412–414
Transimpedance amplifier (TIA), 23, 131, 239, 378–379
Transistor-Outline-can (TO-can), 144–145
Transition metals, 163–164
Transmission electron microscope (TEM), 132
Transmitter (TX), 385
Transmitter dispersion eye closure (TDEC), 239
Transmitter dispersion eye closure quaternary (TDECQ), 93, 239
Transmitter optical subassembly (TOSA), 245–246
Transverse electric polarization (TE polarization), 374–376
Transverse magnetic modes (TM modes), 171, 392–393
Transverse-electric modes (TE0 modes), 171
Traveling wave (TW), 390
Traveling wave MZMs (TWMZMs), 252–253
Traveling wave photodetectors (TWPDs), 128
Tree networks, 5
TSV. *See* Through-silicon vias (TSV)
TSVs. *See* Through silicon/substrate vias (TSVs)
Tunable delay element, 92
Tunable lasers, 38–39, 59, 61–62
  Kerr comb lasers, 40–41
  SGDBR, 39
  Vernier ring lasers, 37–39
Tunable-wavelength lasers, 36
TW. *See* Traveling wave (TW)
TWMZMs. *See* Traveling wave MZMs (TWMZMs)
2D packaging, 214
Two-Photon Grayscale Lithography (2GL®), 427–428
Two-photon polymerization (2PP), 427–428
TWPDs. *See* Traveling wave photodetectors (TWPDs)
TX. *See* Transmitter (TX)
TxRx subsystem requirements, 131–132

U
Unitraveling carrier (UTC), 128
UTC. *See* Unitraveling carrier (UTC)
UV-Ozone, 456

V
V-groove, 52, 394–395
Variation model, 360
VCSELs. *See* Vertical-cavity surface-emitting lasers (VCSELs)
Vector modulator, 203–204
Verification, 371–372
Verilog-A model, 354, 372, 399
Vernier ring lasers, 37–39
Vertical p–i–n PDs, 126–127, 137
Vertical-cavity surface-emitting lasers (VCSELs), 36, 93–94
Voltage-mode approach, 90–91
Volume deployment, development paths for, 219–224
Voxel, 427–428

W
WAC. *See* Wafer acceptance criteria (WAC)
Wafer acceptance criteria (WAC), 464
Wafer bonding, 219
Wafer fab process, 463
Wafer level testing, 321
Wafer qualifications/certification, 463
Wafer-level inline, 376–377
Wafer-level testing, 403–404
Wafer-scale fabrication, 103
Wafer-scale integration of lasers, 60–61
Wall-plug efficiency (WPE), 40
Wave interference, 177
Waveguides, 160–164
  Bragg grating Mux/Demux and filters, 181–183
  comparison of integrated photonic waveguides, 162*t*
  crossings and escalators, 164–167, 166*f*
  implementation of modulators, 74*f*
  manufacturability, 163–164
  waveguide-based modulators, 105–106
  waveguide-integrated PDs, 124–125, 128–129
  waveguide-integrated p–i–n Ge-on-Si photodetector, 125*f*
  waveguide–detector integration paradigm, 124–126
Wavelength
  based on wavelength selection filter, 37*f*
  filters, 160
  multiplexers, 160
  routing switching, 282–283
  tuning, 39
Wavelength blocker (WBL), 286–287
Wavelength division multiplexing (WDM), 160, 175, 217, 282–283, 304–305, 325–326, 335
  architectures, 36
  links, 179–181
  systems, 56–57, 61
Wavelength selecting stage (WSS), 282–283
Wavelength selective switches (WSSs), 295–296
WB. *See* Wire bonding (WB)
WBL. *See* Wavelength blocker (WBL)
WDM. *See* Wavelength division multiplexing (WDM)
Web 2.0, 13–19
  fiber and optical interfaces, 16, 17*t*
  ML hardware, 18–19
  network hardware, 14–15
  workloads, 16–17
Wide-sense and strictly nonblocking networks, 273
Widely tunable lasers (WTL), 461–462
Wire bonding (WB), 214, 418
WPE. *See* Wall-plug efficiency (WPE)
WSS. *See* Wavelength selecting stage (WSS)
WSSs. *See* Wavelength selective switches (WSSs)
WTL. *See* Widely tunable lasers (WTL)

X
X-ray diffraction (XRD), 133–134
X-ray radiography, 447
XRD. *See* X-ray diffraction (XRD)

Y
Y-splitters, 164
Yield analysis and optimization, 361–363

Z
ZB. *See* Zettabyte (ZB)
ZEMAX simulations, 416*f*
Zettabyte (ZB), 265
Zion, 18−19, 18*f*
400ZR deployment, outlook for, 227*f*
400ZR products, 210
ZR + derivatives, 221−223

Integrated Photonics for DATA COMMUNICATION APPLICATIONS

An introduction to the key principles and advances of application-specific photonic integrated circuits (PICs) from leaders across industry and academia.

Integrated Photonics for Data Communication Applications is an important reference that reviews the key concepts, design principles, performance metrics, and manufacturing processes from advanced photonic devices to integrated photonic circuits.

First, there is an overview of the trends and commercial needs of data communication in data centers and high-performance computing, with contributions from end-users presenting the key performance indicators. Then the fundamental building blocks are reviewed—the devices (lasers, modulators, photodetectors, and passive devices) that are the individual elements that make up the photonic circuits. These chapters include an overview of device structure and design principles and their impact on performance. The next section focuses on putting these devices together to design and fabricate application-specific photonic integrated circuits to meet performance requirements. Finally, key areas and challenges critical to the commercial manufacturing of photonic integrated circuits and the supply chains being developed to support innovation and market integration are discussed.

Key features

- Includes contributions from leading experts and end-users across academia and industry working on the most exciting research directions of integrated photonics for data communications applications
- Provides an overview of data communication-specific integrated photonics starting from fundamental building block devices to photonic integrated circuits to manufacturing tools and processes
- Presents key performance metrics, design principles, performance impact of manufacturing variations and operating conditions, as well as pivotal performance benchmarks

About the editors

Madeleine Glick is a senior research scientist in the Lightwave Research Laboratory at Columbia University.

Ling Liao is an Intel Fellow and chief architect of photonic integration in the Silicon Photonics Products Division at Intel.

Katharine Schmidtke is the Director of ASICs and Custom Silicon Sourcing at Meta Platforms.

About the series

The *Integrated Photonics* series is led by Lionel Kimerling, executive at AIM Photonics Institute and Thomas Lord Professor of Materials Science and Engineering at MIT, and Sajan Saini, education director at the Initiative for Knowledge and Innovation at MIT.

Each edited volume features thought leaders from academia and industry in critical application area fronts (including data communications, high-speed wireless, smart sensing, and imaging) and addresses the latest advances in design, fabrication, packaging and testing standards and practices necessary to fulfill near-term systems-level benchmarks.

ELSEVIER elsevier.com/books-and-journals

ISBN 978-0-323-91224-2